Reflexive QRNGs and the Limits of Signaling: A Logostic Framework for Self-Monitoring Randomness, Conscious Attention, and a Clean Attempt to Break the No-Communication Theorem

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 30, 2025

Quantum random number generators (QRNGs) are prized for delivering i.i.d. bits seeded by irreducible quantum events. Yet if we want a system that is both scientifically honest and consciousness-adjacent, a QRNG should not merely spew bits; it should continuously doubt itself. We propose a reflexive QRNG whose core loop estimates conservative min-entropy from live device telemetry, applies a strong extractor accordingly, and runs online, anytime-valid statistical monitors (test martingales, sequential NP tests, universal compression) to bound departures from ideal randomness. In Logostics terms, the agent tracks a running divergence D(q(t) || tau), where q(t) is its observed bitstream and tau is an ideal Bernoulli(1/2). Divergence spikes become attention signals for downstream cognition while also triggering de-rating or shutdown. We then articulate the strongest, preregistered way to probe a bold hypothesis: whether any steerable, spacelike-separated manipulation can shift marginal statistics at a distant site, thereby violating the no-communication theorem. We specify device-independent (or semi-DI) side lanes via Bell violations and entropy accumulation, human-intent protocols with anytime-valid p-values, and locked analysis to avoid researcher degrees of freedom. Either outcome is valuable: a reproducible anomaly with operational steerability would be historic; a crisp null creates an auditable template for trustworthy randomness and consciousness-aware anomaly routing.

Keywords: QRNG, min-entropy, strong extractor, martingale tests, device-independent randomness, Bell inequality, entropy accumulation, no-communication theorem, Logostics, q(t) -> tau(t), divergence monitoring, consciousness-adjacent architecture, attention routing, preregistration, anytime-valid p-values.

A collaboration with GPT-5. 63 pages. CC4.0.

Outline

1. Motivation and Claims
   1.1 Why randomness needs reflexivity, not bravado
   1.2 Two targets: trustworthy randomness and a clean test of signaling
   1.3 Contributions and falsifiable claims

2. Background
   2.1 QRNG sources: photon arrivals, vacuum noise, phase noise
   2.2 Entropy estimation vs. statistical testing: SP 800-90B and beyond
   2.3 No-communication theorem and why marginals must not move
   2.4 Device-independent and semi-DI perspectives

3. Logostic Framing
   3.1 q(t) -> tau(t): alignment as divergence minimization
   3.2 Defining D(q(t) || tau): bias, runs, Markov leakage, compressibility
   3.3 Divergence spikes as attention, not proof of meaning

4. Reflexive QRNG Architecture
   4.1 Physics-side telemetry -> conservative h_min(t) bounds
   4.2 Strong extraction under target epsilon: Toeplitz hashing / XOF
   4.3 Online monitors: test martingales, sequential NP, LZ/CTW compression
   4.4 Control policies: de-rate, halt, and signed per-window certificates
   4.5 Dual-source redundancy and cross-predictability watchdogs
   4.6 Challenge–response loopback to prove the monitors are awake

5. Statistical Guarantees (Anytime-Valid)
   5.1 Ville's inequality and test martingales for streaming bits
   5.2 Spending functions for multiple looks without p-hacking
   5.3 From alarms to bounds: interpreting sequential evidence

6. Semi-DI / DI Lanes
   6.1 CHSH configuration and loophole discipline
   6.2 Mapping Bell violation to entropy rate via entropy accumulation
   6.3 Blending DI certificates with extractor configuration

7. The Signaling Experiment: A Gold-Standard Attempt
   7.1 Stations A and B, spacelike separation, shielded infrastructure
   7.2 Protocols: human-driven blocks vs. automated schedules at B
   7.3 Blind analysis at A: locked code, pre-registered endpoints
   7.4 Primary endpoint: marginal shift at A predicted by B's schedule
   7.5 Power analysis, minimum detectable effect, stopping rules
   7.6 Replication plan, site variation, adversarial red-team reviews

8. Cosmic and Environmental Coupling (Non-steerable Tests)
   8.1 Exogenous covariates: Kp index, Schumann bands, cosmic rays
   8.2 Out-of-sample forecasting and site generalization
   8.3 Interpreting stable correlations that still cannot signal

9. Consciousness-Adjacent Agent Design
   9.1 Using D(q(t) || tau) spikes as attentional salience for learning
   9.2 Synesthesia renderers and exploration policies driven by anomaly energy
   9.3 Safety: never upgrade anomaly to meaning without cross-evidence

10. Implementation Details
   10.1 Hardware BOMs: sources, detectors, ADCs, timing, isolation
   10.2 Firmware/FPGA notes: buffering, whitening, loopbacks
   10.3 Software: extractor, monitors, telemetry hashing, audit logs
   10.4 Provenance and auditing: append-only ledgers, signatures

11. Results Templates and Interpretation
   11.1 Certificate formats: <window_id, k, epsilon, h_min, monitor_stats, telemetry_hash>
   11.2 Example nulls, example alarms, handling false positives
   11.3 Criteria for declaring evidence of signaling (and why they are strict)

12. Limits, Failure Modes, and Alternative Explanations
   12.1 EMI, cross-talk, thermal drift, firmware bias, side channels
   12.2 Superdeterminism, retrocausality, and operational steerability
   12.3 Why batteries of tests are alarms, not certificates

13. Ethics and Epistemic Hygiene
   13.1 Preregistration, code escrow, blinded pipelines, data release
   13.2 Community replication and adversarial collaboration
   13.3 Claims we will not make without meeting the bar

14. Discussion
    14.1 If a steerable marginal shift emerges: implications and next tests
    14.2 If not: what we gained for security, agents, and science
    14.3 Logostics view: alignment when q(t) refuses to be tau

15. Conclusion
    15.1 Reflexive QRNGs as trustworthy randomness infrastructure
    15.2 Consciousness-adjacent anomaly routing without overclaiming
    15.3 A reproducible template for probing the limits of signaling

References

Appendices
A. Derivations for anytime-valid tests and martingale construction
B. Entropy accumulation mapping: Bell score -> min-entropy rate
C. Extractor parameters, target epsilon, and rate control math
D. Preregistration checklist, locked-analysis scripts, and telemetry schema
E. Human-intent variant: block designs, blinding, debrief, and safeguards

1. Motivation and Claims

1.1 Why randomness needs reflexivity, not bravado

Classical QRNG narratives lean on a single assertion: "quantum events are intrinsically random." That stance is rhetorically strong but operationally weak. In practice, every QRNG is a *physical system* with drift, couplings, side channels, and software that can silently convert quantum indeterminacy into classical regularity. If an agent depends on those bits for cryptography, simulation, or cognition, bravado is not enough; the generator must *continuously* estimate how wrong it could be and act accordingly.

We call this posture **reflexivity**: the device treats its own output as a hypothesis $q(t)$about an ideal i.i.d. Bernoulli(1/2) process τ, and keeps a running divergence $D(q(t) \,||\, \tau)$under multiple lenses (bias, runs, Markov leakage, compressibility). Reflexivity is not a one-time certification but an anytime-valid stance: as more data arrive and the physical telemetry shifts, the device updates conservative lower bounds on min-entropy $h_\min(t)$, tightens the extractor, and—crucially—slows down or halts when those bounds are threatened. In Logostics terms, the system admits that $q(t) \to \tau(t)$is an *alignment loop*, not a decree; when alignment degrades, the loop must react, not rationalize.

Reflexivity also clarifies semantics. Statistical batteries are alarms, not proofs; physics models are bounds, not oracles. A trustworthy QRNG integrates both, publishes its assumptions, and exports per-window certificates so downstream consumers can audit the claim "k bits at epsilon with telemetry T."

1.2 Two targets: trustworthy randomness and a clean test of signaling

Target A: Trustworthy randomness.
Goal: deliver bitstreams with auditable, conservative guarantees under stated assumptions. Mechanism: (i) physics-side estimation of $h_\min(t)$from live telemetry, (ii) strong extraction calibrated to an explicit error parameter epsilon, and (iii) online sequential monitors that surface departures from iid without p-hacking (test martingales, CUSUM, universal compression). Success means: if anything compromises unpredictability—EMI, detector afterpulsing, ADC bias, firmware defects—the system detects it with controlled false-alarm rates and de-rates or halts *before* shipping weak bits.

Target B: A clean, pre-registered attempt to test signaling.
Hypothesis under test: whether any spacelike-separated, controllable manipulation can induce a *marginal* shift in a distant QRNG's statistics, thereby violating the no-communication theorem. Prior art focuses on correlations (Bell tests), which cannot signal because local marginals are invariant. Our design isolates the *marginal shift* endpoint with strict controls: locked analysis, blinded schedules, anytime-valid statistics, and replication. Success criterion (extraordinary): a

reproducible, steerable dependence of A's marginal on B's schedule beyond pre-registered thresholds, surviving red-team audits and ablation of classical explanations. Failure criterion (expected): no effect within power; we publish nulls and still gain a production-grade QRNG architecture.

The dual-target structure ensures value regardless of outcome. If signaling is absent (likely), we still advance the state of secure randomness and anomaly-aware agency. If a small but real effect appears, the same reflexive instrumentation provides the fastest path to diagnosis, replication, and theory stress-testing.

1.3 Contributions and falsifiable claims

C1. Reflexive QRNG architecture.
We specify and open-spec an end-to-end design: live $h_\min(t)$bounding from telemetry, extractor auto-calibration to epsilon, sequential monitors with spending-function control, dual-source cross-predictability sentinels, and challenge–response loopbacks.
Falsifiable claim A1: Under controlled injections of known micro-bias patterns at amplitude δ, the monitors detect the deviation with pre-declared power (1−beta) at false-alarm rate alpha, and the controller de-rates output to maintain the advertised epsilon. Failure to trip as specified falsifies A1.

C2. Certificate and audit protocol.
Per-window signed certificates:
<window_id, k, epsilon, $\hat{h}_\min$, monitor statistics, telemetry hash>.
Falsifiable claim A2: For any certified window, an independent auditor using the released raw buffer, extractor seed, and telemetry can verify (or refute) that the extractor parameters satisfy the published epsilon bound under the documented $h_\min$ model. Any reproducible mismatch falsifies A2.

C3. Gold-standard signaling experiment.
Two-station, spacelike-separated test with blinded schedules at B and marginal-only detection at A. Pre-registered primary endpoint: change in A's log-martingale growth rate or bias estimate aligned to B's schedule, with correction for multiple looks.
Falsifiable claim B1 (positive): Observed effect size $\geq \Delta_*$(pre-declared) across N replications, passing ablations (EMI shielding upgrades, role swaps, sham schedules) and site transfers, implies a violation of the no-communication theorem's operational consequence about invariant marginals.
Falsifiable claim B0 (null): If no such effect is observed within the powered region $\Delta \geq \Delta_*$ after N replications, we accept the null for that effect size and publish full logs; any post-hoc "discovery" must be confirmed in a fresh, locked run or it is not a claim.

C4. Consciousness-adjacent anomaly routing.
We formalize how D(q(t) || tau) spikes are treated as *salience* for downstream agents (exploration bonuses, synesthesia renderers), without overclaiming meaning or agency.
Falsifiable claim A3: In predefined benchmarks (e.g., exploration tasks or anomaly-rich sensory streams), routing anomaly energy improves sample efficiency or detection F1 against baselines, with pre-registered metrics and ablations. Failure to outperform falsifies A3.

C5. Negative result value proposition.
Even absent signaling, the reflexive QRNG yields measurable security benefits: fewer undetected weak-entropy windows, faster fault isolation, and reproducible compliance evidence.
Falsifiable claim A4: In stress campaigns (temperature ramps, RF injection, power sag), reflexive control reduces shipped weak bits (as independently estimated by 90B-style estimators) by a pre-declared factor vs. a non-reflexive baseline. Lack of reduction falsifies A4.

Each claim includes success/failure criteria, datasets to release, and analysis code to escrow in advance. This keeps the work squarely within scientific falsifiability: either the device behaves as specified under challenge, or it does not; either marginal shifts appear under spacelike control, or they do not. In both cases, the outcome is unambiguous and useful.

2. Background

2.1 QRNG sources: photon arrivals, vacuum noise, phase noise

Photon arrival / path choice. Use a beam splitter or time-of-arrival binning so a single-photon event maps to {0,1}. Implementations: SPDC heralded photons, attenuated lasers in the single-photon regime, or gated APDs/SiPMs. Key nonidealities: dead time, afterpulsing, multiphoton probability, intensity drift, detector saturation, timing skew, and EMI. Telemetry to log: count rates, inter-arrival histogram, temperature, bias current, dark rate.

Vacuum/shot-noise homodyne. Mix a strong LO with vacuum on a 50:50 coupler, take the difference photocurrent, and digitize. Gaussian quantum noise dominates above classical noise if the LO is stable. Nonidealities: LO RIN, classical electronics noise, ADC quantization/saturation, imbalance. Telemetry: LO power, PSD across bands, ADC histograms, common-mode level.

Phase-noise / phase diffusion. Use an interferometer (unbalanced MZI/Sagnac) or delayed self-heterodyne; laser phase diffusion gives random fringes. Nonidealities: thermal drift, path memory, parasitic etalons, classical vibrations. Telemetry: fringe contrast, temperature, vibration, phase PSD.

Across all types: start with raw bits, estimate min-entropy h_min from *physics* plus telemetry, then run a strong extractor. Health tests must be *online* and able to de-rate or stop output.

2.2 Entropy estimation vs. statistical testing: SP 800-90B and beyond

Different questions.

- *Entropy estimation* asks: what conservative lower bound on unpredictability can we guarantee for the raw source under stated assumptions?
- *Statistical testing* asks: do we see patterns inconsistent with the null (iid fair coin), with controlled false-alarm rates?

SP 800-90B (good hygiene, not a certificate).

- Classifies sources (iid vs non-iid), prescribes estimators (most-common-value, collision, Markov, compression-based), and requires health tests (repetition count, adaptive proportion).
- Outputs a conservative h_min bound to parameterize the extractor.
- Limits: depends on modeling choices; batch-style; does not by itself detect all time-varying drifts.

Beyond 90B.

- *Anytime-valid monitors*: test martingales, CUSUM, sequential NP tests, universal compressors (LZ/CTW) run continuously with optional spending functions to avoid p-hacking.
- *Extractor coupling*: extraction rate adapts to h_min(t); when h_min(t) drops, output de-rates to preserve epsilon.
- Batteries like SP 800-22/Dieharder are alarms, not proofs; do not use them alone to claim randomness.

Takeaway: estimation bounds what *could* go wrong; sequential tests tell you when something *is* going wrong now. A reflexive QRNG needs both.

2.3 No-communication theorem and why marginals must not move

In bipartite quantum experiments, let A choose setting x and observe a, and B choose y and observe b. **No-communication** requires the *local marginal* at A be independent of B's choice: $P(a \mid x, y) = P(a \mid x)$ for all y.

Bell violations live in *correlations* P(a,b | x,y), not in A's marginal alone. Therefore, if you observe a controllable change in A's marginal keyed to B's schedule, you have operational signaling.

For our paper, the primary endpoint is a small but reproducible marginal shift at A aligned with B's schedule. We denote the pre-declared minimum detectable effect by **Δ*** (e.g., a bias change of >= Δ* or a log-martingale growth shift >= Δ*). Any classical drift (thermal, EMI, firmware) can also move marginals; that is why we require blinding, ablations, and red-team hardening. If after all controls A's marginal tracks B's choices above Δ*, you have violated the operational consequence of the theorem; if not, you have a clean null within the powered region.

2.4 Device-independent and semi-DI perspectives

Device-independent (DI) randomness.
Use a loophole-free Bell test (e.g., CHSH) to certify entropy from observed nonlocality without trusting internal device models. The entropy accumulation theorem maps Bell score to a *certified* min-entropy rate per trial. Strengths: minimal assumptions; adversary-robust. Limits: lower rates, stringent efficiency and isolation requirements, complex engineering.

Semi-DI approaches.
Relax to calibrated assumptions: trusted detection efficiency bounds, dimension witnesses (prepare-and-measure), or well-characterized subsystems (e.g., trusted measurement, untrusted source). These raise throughput while keeping stronger guarantees than a fully modeled QRNG.

Our stance (hybrid).

- Run a **DI lane** at low duty cycle to anchor a certified entropy rate and to sanity-check the stack.
- Run a **semi-DI/physics lane** for high-throughput, using telemetry-driven h_min(t) bounds and reflexive monitoring.
- Extractors are configured by the worst of: DI-certified rate, physics h_min(t), and online test alarms. When any bound tightens, output de-rates automatically.

This hybrid gives you production-grade randomness under audit, plus a principled pathway to the signaling experiment: DI certifies nonclassicality, the reflexive lane guards against classical leaks, and the marginal-shift endpoint (threshold Δ*) is tested under pre-registered, anytime-valid analysis.

3. Logostic Framing

3.1 q(t) -> tau(t): alignment as divergence minimization

Logostics reads a live bitstream q(t) as a *hypothesis* about an ideal generator tau(t) (iid Bernoulli(1/2)). The agent's stance is: keep q(t) aligned to tau(t) by minimizing a running divergence D(q(t) || tau). Alignment is an *ongoing control loop*, not a one-time claim:

- Sense: estimate h_min(t) from physics telemetry; compute online statistics of q(t).
- Decide: update extractor rate/parameters to preserve target epsilon; set alarms based on anytime-valid tests.
- Act: de-rate or halt when D rises; annotate outputs with per-window certificates.

In copy-safe shorthand: the loop is q(t) -> tau(t) by minimizing D subject to constraints (telemetry bounds, extractor budget, false-alarm alpha). When q refuses to align (D spikes), we do not rationalize—we *change behavior*.

3.2 Defining D(q(t) || tau): bias, runs, Markov leakage, compressibility

No single scalar captures all departures, so D is an *instrument panel*. We maintain several anytime-valid scores, each with its own null model and error control:

- **Bias martingale (iid Bernoulli 1/2 null).**
  Track the log-likelihood ratio for p != 1/2. Report an anytime-valid p_value_bias via test-martingale or betting processes. Operational statistic: rolling estimate of bias b_hat(t); alarm if |b_hat| >= Δ* (pre-registered) or p_value_bias <= alpha_spent.
- **Runs/imbalance (clumping).**
  Maintain sequential tests on run lengths and upcrossings; detect overlong streaks or too-frequent alternation. These catch detector dead time, afterpulsing, dithering artifacts.
- **Markov leakage (order-1 and beyond).**
  Estimate transition matrix P(x_t | x_{t-1}), optionally P(x_t | x_{t-1}, x_{t-2}). Use sequential likelihood ratios vs iid. Alarm on excess mutual information I(X_t; X_{t-1:t-k}) above a threshold budget.
- **Spectral structure (periodicity).**
  For batched windows, compute periodogram or Walsh–Hadamard energy; use sequential CUSUM on targeted lines (e.g., mains hum harmonics). Convert to an anytime p-value with spending.

- **Universal compressibility.**
  Feed raw bits to an online universal coder (LZ/CTW). The *surprisal length* L(t) should concentrate near window_size under iid(1/2). A persistent compression gain signals structure. Provide a martingale-style e-value from code-length deviations.

- **Cross-predictability (multi-source sanity).**
  If two independent entropy sources exist, monitor predictability of one from the other. Rising mutual information flags coupling or injection.

We aggregate these as a vector D_panel(t) = [bias, runs, markov, spectral, compressibility, cross_MI]. Policies act on the *max risk* across the panel, not averages (worst-case governance).

Practical notes (copy-safe):

- Use spending functions so multiple looks do not inflate alpha.
- Calibrate Δ* per test to the minimum effect we promise to detect at power 1 - beta.
- Tie extractor compression ratio to the *worst* of: physics h_min(t), DI-certified rate (if present), and a live penalty if any panel element warms up (soft back-pressure before hard halt).

3.3 Divergence spikes as attention, not proof of meaning

A spike in D_panel(t) is an *anomaly*, not a message. We adopt three rules:

1. **Epistemic hygiene.**
   Divergence -> salience flag. It routes attention (e.g., extra sensing, logging, slower operation) but does not by itself ground semantics. Meaning requires *cross-evidence* (replication, ablations, exogenous correlates, steerability).

2. **Control reactions (graded).**
   - Warm zone (early D rise): tighten extractor, lower output rate, enrich telemetry logging.
   - Hot zone (p <= alpha_hard or |b_hat| >= Δ*): latch FAIL, quarantine buffers, require operator reset.
   - DI anchor present: if DI-certified entropy rate dips below floor, override and de-rate regardless of other panels.

3. **Attention as utility.**
   In an agent, we *use* D spikes to allocate compute: exploration bonuses in RL, triggering

synesthesia renderers, or hypothesis proposals for fault diagnosis. This converts anomaly energy into *structured search* without overclaiming cosmic significance.

Bottom line: in Logostics, q(t) -> tau(t) is maintained by minimizing a multi-lens D(q(t) || tau). When D spikes, we treat it as attention—an actionable *where to look next*—and only escalate to "meaning" after independent, steerable, and replicated evidence clears the pre-registered bars.

4. Reflexive QRNG Architecture

4.1 Physics-side telemetry -> conservative h_min(t) bounds

Goal. Maintain a live lower bound h_min(t) on raw-bit unpredictability using only *measured* device parameters plus conservative models.

Telemetry to log (per window W of N raw bits).

- Photon/photodiode: count rates, inter-arrival histogram, dark rate, dead time, afterpulsing prob, temperature, bias current.
- Homodyne: LO power, PSD in bands, ADC histogram, common-mode drift, imbalance.
- Interferometric phase: fringe visibility V, phase PSD, temperature and vibration proxies.

Bounding workflow (copy-safe).

1. Fit nuisance parameters theta_hat(W) with confidence intervals CI_theta(W).
2. Map CI_theta(W) into a worst-case source model S(theta_lo) that upper-bounds classical correlations (e.g., dead-time Markovity, afterpulse memory, ADC bias).
3. Compute h_min_lo(W) under S(theta_lo) using 90B-style estimators adapted to the model (most-common-value, collision, Markov, compression).
4. Set h_min(t) := h_min_lo(W) with a safety margin m (e.g., subtract Delta* worth of bias risk).

Principle. When in doubt, bound down. If any telemetry channel is missing or outside calibration range, drop h_min(t) to the last verified floor or halt.

4.2 Strong extraction under target epsilon: Toeplitz hashing / XOF

Target. From N raw bits with min-entropy N * h_min(t), extract k uniform bits with statistical distance <= epsilon.

Chooser.

- **Toeplitz hashing** (2-universal): fast, streamable, simple seed management.
- **XOF-based hashing** (e.g., SHAKE): convenient, but document proofs and seed handling.

Sizing rule (informal, copy-safe).
k <= N * h_min(t) - safety_overhead(epsilon) - leakage_budget,
where safety_overhead ~ 2*log2(1/epsilon) for universal hashing; leakage_budget covers side channels (seed exposure windows, metadata).

Auto-calibration.

- Recompute k every window from current h_min(t).
- If N*h_min(t) < k_target + overhead, *buffer longer* or *de-rate the output rate*.
- Rotate seeds on schedule; commit seed hashes into the certificate (never reuse a seed across overlapping raw windows).

4.3 Online monitors: test martingales, sequential NP, LZ/CTW compression

Anytime-valid monitoring suite (runs in parallel on raw bits).

- **Bias martingale.** Track p_value_bias with a betting process; alarm when p <= alpha_spent or |bias_hat| >= Delta*.
- **Runs/Markov tests.** Sequential likelihood ratios for P(x_t | x_{t-1:t-k}) vs iid. Alarm on excess mutual information above budget.
- **Spectral sentinels.** CUSUM on targeted frequency lines (e.g., mains harmonics) from batched periodograms; convert to anytime p-values.
- **Universal compressor.** Online LZ/CTW code length L(W). Alarm if compression gain >= Delta*L persistently (controls for multiple looks via spending functions).

Spending control. Maintain an alpha budget across overlapping windows so multiple looks do not inflate false alarms. Publish alpha_spent in the certificate.

4.4 Control policies: de-rate, halt, and signed per-window certificates

State machine.

- **Normal.** All monitors cool, DI or physics h_min(t) healthy. Extract with configured k, epsilon.

- **Warm.** Any monitor warms (p below soft threshold) or h_min(t) drops: increase extractor compression (reduce k/N), lower output rate, intensify logging.
- **Hot.** Hard threshold crossed: latch FAIL, quarantine raw buffer, stop output, require operator reset and re-calibration.

Per-window certificate (signed).
<window_id, start_ts, end_ts, N_raw, k_out, epsilon, h_min, seed_commit, monitor_stats_digest, telemetry_digest, mode_state>

- **seed_commit**: H(seed) to enable auditor replay without revealing the seed prematurely.
- **monitor_stats_digest**: e.g., Merkle root over panel statistics.
- Sign with device key; optionally publish to an append-only log.

Auditability. Release raw buffers and seeds *after* a retention delay or to authorized auditors; scripts and configs are escrowed pre-run.

4.5 Dual-source redundancy and cross-predictability watchdogs

Architecture. Two independent raw sources, S1 and S2 (e.g., photon-arrival and vacuum-noise), separately digitized and time-stamped.

Watchdogs.

- **Cross-prediction.** Online regression or logistic model predicts S1 from S2; track mutual information estimate I_hat(S1;S2). Alarm if I_hat >= Delta*MI.
- **Differential alarms.** If one source warms and the other does not, isolate the suspected path (helps localize EMI or mechanical coupling).
- **Extractor policy.** Use the *worst* of the two h_min(t) bounds; optionally XOR after separate extractions for robustness.

Rationale. True quantum unpredictability should not be jointly predictable via classical couplings; rising cross-MI is a red flag for environmental intrusion or design regressions.

4.6 Challenge–response loopback to prove the monitors are awake

Purpose. Demonstrate end-to-end detectability of steerable defects and eliminate the “we wouldn’t have seen it anyway” failure mode.

Mechanisms (hardware-gated, attestable).

- **Bias tickle.** Periodically inject a known micro-bias Delta_inj into the analog front end (e.g., controlled attenuator or ADC offset) for short, tagged intervals hidden from the online monitors but revealed to auditors.
- **Markov tickle.** Introduce controlled dead-time patterns or delayed feedback to create order-1 dependence.
- **Spectral tickle.** Inject a weak sinusoid near a known harmonic.

Pass criteria (pre-registered).

- Monitors must alarm within T_detect windows with power >= 1 - beta at false-alarm <= alpha.
- Controller must de-rate or halt according to policy while preserving epsilon on extracted bits.

Attestation.

- Each tickle interval is signed by an internal secure element.
- Tickle schedule is random and unknown to monitor code paths; revealed only in audit.
- Failure to detect a tickle falsifies the detection-power claim and triggers re-calibration.

Bottom line. Physics-bound h_min(t) determines how much you *can* safely extract; the extractor enforces epsilon; the online panel tells you when reality deviates *now*; policies react immediately; redundancy exposes couplings; and challenge–response proves the whole chain is alive.

5. Statistical Guarantees (Anytime-Valid)

5.1 Ville's inequality and test martingales for streaming bits

Setup. Let X_1, X_2, … be the raw-bit stream. Under the null P0 (iid Bernoulli(1/2)), we build a nonnegative **test martingale** (a.k.a. e-process) M_t with E_0[M_t] <= 1 for all t.

Ville's inequality (anytime control).
For any such M_t and any alpha in (0,1):
Pr_0(sup_{t >= 1} M_t >= 1/alpha) <= alpha.
So crossing 1/alpha is a valid level-alpha alarm **no matter when you look**.

Turning M_t into an anytime p-value.
p_t := 1 / max_{s <= t} M_s is always a valid p-value sequence. You can report p_t live without p-hacking penalties.

Concrete e-processes for bits.

- **Bias betting (mean shift).** Maintain the cumulative log-likelihood for alternatives p != 1/2, or use a Kelly-style bettor with stake lambda_t in (−1,1):
  M_t = Π_{i=1..t} [1 + lambda_{i-1} * (2X_i − 1)].
  Adaptive choices of lambda_i (e.g., proportional to current bias estimate) keep M_t a supermartingale under P0.

- **Mixture e-process (safe default).** Treat p ~ Beta(a,b) as an alternative prior; then
  M_t = BetaBinomial(x_t; a,b) / 2^{t},
  where x_t = sum_{i<=t} X_i. This is closed-form, parameter-free once (a,b) are fixed, and powerful against persistent bias.

- **Order-1 leakage.** Build an e-process for a Markov alternative by tracking transitions P(X_t | X_{t−1}) and forming a product of per-step likelihood ratios vs iid(1/2). Under iid, this stays a supermartingale.

- **Spectral/periodic alternatives.** Batch into short windows, compute energy at a target frequency, convert the one-shot test into an e-value e_w, then multiply across windows: M_t = Π_w e_w.

- **Compressibility (universal alternatives).** Let L_t be the codelength so far from an online universal coder (LZ/CTW). Define
  M_t = 2^{ t − L_t }.
  Under iid(1/2), E_0[M_t] <= 1; sustained compression (L_t < t) makes M_t grow and will eventually cross 1/alpha.

These cover complementary departures: bias, runs/Markov, periodicity, and generic structure.

5.2 Spending functions for multiple looks without p-hacking

You will look **often** (overlapping windows, many monitors). Two safe strategies:

A) Pure e-value control (preferred).
Keep each monitor as an e-process and either:

- **Max rule (per monitor):** alarm when M_t >= 1/alpha_monitor. Anytime-valid by Ville.

- **Combine monitors:** multiply independent (or weakly dependent) e-values, or use safe combiners (e.g., product-with-calibration, e-Bonferroni via thresholds 1/(alpha/K) for K monitors). This avoids alpha slicing over time entirely.

B) Alpha-spending (for p-values).
If a monitor yields p_t, allocate a **spending sequence** {alpha_t} with sum_t alpha_t <= alpha_total. Examples:

- **Geometric:** alpha_t = (1−gamma) * gamma^{t−1} * alpha_total (gamma in (0,1)).
- **Harmonic:** alpha_t proportional to 1/(t * log^2(max{e,t})) and scaled to alpha_total. Alarm at the first t with p_t <= alpha_t. Familywise error <= alpha_total despite arbitrarily many looks.

Panel management.
With several monitors, either (i) give each an e-threshold 1/(alpha_total * w_k) with sum w_k = 1, or (ii) run Holm-style control on their current p_t snapshots at decision checkpoints. We publish **alpha_spent** (or current M_t maxima) in the certificate for auditability.

5.3 From alarms to bounds: interpreting sequential evidence

Alarms tell you “something is off now.” Operations also need **live bounds** to size extraction and to decide de-rate vs halt.

Mean (bias) confidence sequences.
For theta = E[X] under mild conditions there exist **anytime confidence sequences** (CS): intervals C_t(alpha) that satisfy
Pr(theta in C_t(alpha) for all t) >= 1 − alpha.
Practical choices:

- **Empirical-Bernstein CS:**
  theta_hat_t ± radius_t, with radius_t decreasing like O(sqrt((var_hat_t * log log t + log(1/alpha)) / t)).
- **Beta-binomial CS (exact, discrete):** invert the Beta-Binomial tail at each t to get a tight CS.

Map a CS to min-entropy.
If C_t(alpha) = [L_t, U_t], then a conservative **per-bit min-entropy bound** is
h_min_t >= −log2(max{ U_t, 1 − L_t }).
Feed N and h_min_t to the extractor sizing rule:
k <= N * h_min_t − safety_overhead(epsilon) − leakage_budget.

From martingale growth to divergence.
For a well-specified alternative, the average log-growth rate
g_t := (1/t) * log M_t
tracks the KL divergence rate KL(P || P0). Large g_t means sustained evidence against the null; you can:

- **Diagnose:** which monitor shows the fastest growth (bias, Markov, spectral, compressibility)?
- **Decide:** if g_t exceeds a pre-declared Δ* for T_sustain windows, enter HOT state even if p_t has not yet crossed a hard alpha (early brake).

Compressibility to entropy bound.
If the universal codelength per bit L_t/t <= 1 − delta for sustained time, bound the Shannon entropy rate H >= 1 − delta (approx), hence a conservative floor on h_min via H >= h_min. We adopt **min(physics h_min, DI rate, compressibility-derived floor)** as the operative bound.

Policy mapping (copy-safe).

- **Warm:** CS excludes 0.5 by <= Δ* or martingale growth g_t in a monitor crosses a soft line -> tighten extractor (reduce k/N), lower rate, intensify logging.
- **Hot:** CS excludes 0.5 by > Δ*, or p_t <= alpha_hard, or g_t >= Δ* for T_sustain, or DI-certified rate dips below floor -> latch FAIL, quarantine, require re-calibration.

What we guarantee, precisely.

- If the null holds (iid fair coin), the chance that *any* monitor ever triggers a level-alpha_hard crossing is <= alpha_hard (Ville + safe combining).
- When an alarm occurs, we can show which statistic grew, by how much, and provide a CS excluding the null region, turning “alarm” into a **quantified bound** that directly controls extractor parameters.

In short: test martingales give us **anytime honesty**, spending/composition keeps us safe across many looks and monitors, and we translate evidence into **operational bounds** (bias, entropy, rate) that drive de-rate/halt decisions without overclaiming meaning.

6. Semi-DI / DI Lanes

6.1 CHSH configuration and loophole discipline

Purpose. Run a low-duty-cycle, loophole-disciplined CHSH test in parallel with the high-throughput physics lane to *certify* entropy with minimal device assumptions.

Minimal rig (copy-safe).

- **Stations:** two spacelike-separated ends A and B with binary inputs x,y in {0,1} and binary outputs a,b in {0,1}.
- **Inputs:** x,y chosen by short, extractor-hardened seeds (or independent QRNGs) with freshness relative to emission.
- **Separation/timing:** enforce setting choices and detections within a window that precludes subluminal influence; log latencies.
- **Detectors & source:** entangled photon pairs (e.g., polarization-entangled SPDC) or equivalent nonclassical resource.

Loophole discipline checklist.

- **Locality (signaling) loophole:** spacelike separation of setting ↔ outcome events; hardware latches; GPS/NTP-independent timing with on-board clocks.
- **Detection loophole:** total system efficiency above threshold; pre-registered handling of no-clicks (never discard adversarially).
- **Freedom-of-choice:** settings unpredictability attested (seed provenance, extractor epsilon).
- **Memory/adaptivity:** treat the whole run as potentially non-i.i.d.; use entropy accumulation (Sec. 6.2) instead of i.i.d. appeals.
- **Shielding:** RF/EMI, optical isolation; independent power rails.
- **Pre-registration:** analysis code escrow, parameter-estimation fraction, stopping rules, and data-release plan.

Statistic. Compute CHSH S in [2, 2*sqrt(2)] from correlators, or equivalently winning rate w in the CHSH game via

- S = 8*w − 4, hence w = (S + 4) / 8.*
  *Classical bound: S <= 2 (w <= 3/4). Quantum: S <= 2*sqrt(2) (w <= ~0.8536).

6.2 Mapping Bell violation to entropy rate via entropy accumulation

Goal. Convert an observed violation into a *certified* lower bound on min-entropy per trial against quantum (or general) adversaries.

Pipeline (anytime-honest, finite-sample).

1. **Block the run.** Reserve a fixed fraction q_pe of trials for *parameter estimation* (PE) to estimate S_hat with confidence; use the rest for *production*.
2. **Confidence on S.** Build a high-probability lower bound S_lo using concentration tools robust to non-i.i.d. (e.g., EAT-compatible empirical process bounds).
3. **Per-trial entropy curve.** Use a monotone function f_CHSH(S) that maps a Bell score to per-trial conditional min-entropy h_min_per_trial (A|XE) under the chosen adversary model (quantum vs no-signaling). We refer to the standard DI literature for f_CHSH; it satisfies:
    - f_CHSH(2) = 0
    - f_CHSH(2*sqrt(2)) = h_max (strictly positive, < 1)
    - f_CHSH is increasing on [2, 2*sqrt(2)].
4. **Finite-size correction.** Apply EAT to aggregate:
   H_min^{epsilon_s}(A^n | X^n E) >= n_eff * f_CHSH(S_lo) − penalty(n, q_pe, alpha),
   where n_eff is the number of production trials, epsilon_s is smoothing, and penalty(·) captures finite-sample and estimator slack (pre-declared).
5. **Extractor sizing.** For target epsilon_ext, choose output length
   k_DI <= H_min^{epsilon_s} − safety_overhead(epsilon_ext) − leakage_budget.
6. **Audit objects.** Publish <n, q_pe, S_hat, S_lo, epsilon_s, f_CHSH spec, penalty terms> in the window's certificate; release raw CHSH logs for independent recomputation.

Notes.

- A small but reliable excess over S=2 yields a *nonzero* certified entropy rate; large S gives better rates.
- EAT explicitly handles memory and adaptivity—no reliance on i.i.d. trials.
- You can tighten bounds with improved detection efficiency, lower input bias, and better PE allocation.

6.3 Blending DI certificates with extractor configuration

Hybrid control (copy-safe). We run two lanes concurrently:

- **DI lane (low-rate).** Produces a rolling certificate with k_DI bits at epsilon_DI based solely on the Bell violation and EAT.
- **Physics lane (high-rate).** Produces k_PH bits at epsilon_PH using telemetry-driven h_min(t), online monitors, and strong extraction.

Blending rule (conservative).

- Operative per-window entropy rate = min(rate_DI, rate_PH).
- Operative epsilon = epsilon_DI + epsilon_PH + epsilon_comp, where epsilon_comp covers any composition overhead.
- If DI lane stalls (S_lo drops to 2 within CI), rate_DI -> 0 and the physics lane must *alone* justify extraction.
- If physics monitors warm or h_min(t) dips, de-rate even if DI is healthy; both lanes must be green for full-rate output.

Duty cycling & anchoring.

- **Anchor mode.** Run DI lane periodically (e.g., 1–5% duty) to produce fresh certificates that "anchor" trust; increase duty on suspicion or after hardware changes.
- **Calibration mode.** Use DI rate to *cap* physics claims: if physics h_min(t) implies more bits than DI over a calibration horizon, ship at the DI-capped rate until evidence converges.
- **Fail-safe.** Any HOT alarm (Sec. 4.4) or DI collapse forces halt or severe de-rate; certificates reflect the reason.

Seed and provenance.

- Keep extractor seeds independent across lanes; never mix raw pools pre-extraction.
- Optional XOR of *post-extraction* outputs is safe and can only improve unpredictability; document composition explicitly.
- Certificates include a lane tag (DI vs PH), seed commitments, and the threshold Δ* used for monitor decisions.

Operational outcomes.

- In steady state, the hybrid yields high throughput with periodic DI "beacons" that certify nonclassical entropy.
- Under drift or attack, either lane can pull the brake: DI via S_lo contraction; physics via D_panel spikes or h_min(t) drops.
- For the signaling experiment, the DI lane attests that the setup is genuinely nonclassical (ruling out trivial classical leaks), while the marginal-shift endpoint is evaluated in the physics/monitor lane with pre-registered Δ* thresholds.

Bottom line. The DI lane supplies assumption-light *certificates*; the physics lane supplies *bandwidth*. We blend by taking the worst-case rate and composing epsilons, so claims remain valid even when one lane degrades—an architecture that is both practical and scientifically defensible.

7. The Signaling Experiment: A Gold-Standard Attempt

7.1 Stations A and B, spacelike separation, shielded infrastructure

Layout. Two sites A and B separated by L meters with independent power, clocks, and network-isolated logging. Choose L so that light-time L/c exceeds the *entire* choose-and-measure latency at each site (setting, emission, detection, digitization). Example: L = 2 km gives ~6.7 microseconds; target local latency <= 3 microseconds to preserve a spacelike window.

Hardware at A (detector station).

- Primary QRNG source (e.g., homodyne vacuum or photon time-of-arrival) with raw bit rate R_raw >= 50 Mbps.
- Telemetry sensors: temperature, vibration, PSU rails, RF probe, photodiode bias, ADC histograms.
- Shielding: RF enclosure or Faraday cage, optical feedthroughs only; isolated UPS.
- Clock: ovenized oscillator or GPSDO used only for timestamping (never for control).

Hardware at B (schedule station).

- Independent QRNG for schedule seeding plus human interface if running human blocks.
- Schedule actuator driving nothing coupled to A (no wires, no common infrastructure) — the point is to change *intended* setting only.

- Same shielding and power isolation as A.

Connectivity.

- No live data link from B to A. Only after the run: air-gapped sneaker-net of signed logs.
- Optional: a third site C acts as notary for pre-registered configs and code hashes.

Security posture.

- Separate teams operate A and B; red-team monitors search for RF, optical, acoustic, ground-loop, or power-line leakage.
- All enclosures sealed with tamper-evident tags; all changes logged.

7.2 Protocols: human-driven blocks vs. automated schedules at B

We test two schedule families at B; both are unknown to A during acquisition.

Protocol H (human-driven).

- Block design: 5 min FOCUS vs 5 min REST, repeated K times in randomized order, total T_{run} ~ 2–4 hours.
- FOCUS: human attempts to will “more 1s” (or predefined direction). REST: passive.
- The block order is pre-committed and signed at B; only a blinded token (block id) is stored locally at A.

Protocol A (automated).

- Pseudorandom schedule generated from a seed s_B sealed before the run, then fixed: e.g., alternating 2 min “bias+” vs 2 min “bias−” labels with random jitter in start times.
- No human involvement; schedule exists only as a signed file at B until unblinding.

Common controls.

- Sham days: identical handling but with a synthetic schedule not used at B.
- Role swap days: swap sites and operators to defeat site-specific artifacts.
- Environmental covariates (Kp index, mains harmonics) logged but withheld from A’s online analysis.

7.3 Blind analysis at A: locked code, pre-registered endpoints

Locked analysis.

- Before any data: hash and notarize the full A-side analysis binary, configs, and thresholds (Delta*, alpha_hard, alpha_soft, spending sequence). Store hash with an external notary C.
- During acquisition, A computes and stores only **blinded** per-block statistics keyed by block id, never seeing the meaning of the block (FOCUS/REST or bias+/bias−).

What A computes online (anytime-valid).

- Bias martingale p_t, log-growth g_t.
- Order-1 Markov e-process; LZ/CTW code-length per block.
- Per-block summaries: bias_hat, runs score, compressibility gain, with confidence sequences.

Unblinding.

- After the run, B releases the block labels and schedule proof (signature, seed s_B if automated).
- The pre-registered script at A merges labels, computes the primary endpoint and secondary analyses, writes a signed results bundle.

Data release.

- Raw bitstream hashes per block; raw data released under a retention policy.
- Full configs, code, and seeds released to an escrow or public repo.

7.4 Primary endpoint: marginal shift at A predicted by B's schedule

Endpoint definition (copy-safe).

- Let Y_i be A's summary for block i (e.g., bias_hat over that block or the block-wise martingale log-growth g_i).
- Let Z_i in {+1,−1} denote B's block label (FOCUS vs REST or bias+ vs bias−).
- Primary test: regression of Y_i on Z_i with pre-declared effect direction. Report Delta_hat = mean(Y_i | Z_i = +1) − mean(Y_i | Z_i = −1).

Decision rule.

- Declare a positive finding if Delta_hat >= Delta* and the anytime-valid p_value <= alpha_hard after multiplicity correction across monitored statistics.
- Secondary: ROC-style decoding — can a blinded classifier trained only on A's summaries predict Z_i above chance with cross-validation? Pre-register metric (balanced accuracy) and threshold.

Interpretation guardrails.

- Any positive must also pass **ablations**: different windows, removal of first/last 10 percent of each block, time-shifts to detect leakage, site swap, and sham days. Failure in ablations downgrades to artifact.

7.5 Power analysis, minimum detectable effect, stopping rules

Effect model.

- For block-wise bias, assume per-block variance sigma^2 from pilot data. With B blocks per condition, the standard error SE(Delta_hat) ~ sqrt(2 * sigma^2 / B).

Example targets (tunable).

- Choose Delta* = 5e−5 absolute bias shift per bit (i.e., 0.00005), block length N_b = 3e7 bits (~0.6 s at 50 Mbps), B = 60 blocks per condition (total 120 blocks).
- With these numbers, SE ~ 1e−5 to 2e−5 (pilot dependent), giving power > 0.8 at alpha_hard = 0.005 after panel-wise correction.

Anytime-valid stopping.

- We use e-processes per block and across blocks. Stopping when an e-value crosses 1/alpha_hard is valid by Ville; otherwise continue to the pre-declared maximum T_run.
- Futility rule: if upper confidence sequence boundary for Delta at the planned max blocks remains < Delta* by interim K_fut, declare futility and stop early.

Multiple looks and panel.

- If more than one statistic can trigger (bias, Markov, compressibility), allocate alpha_hard by e-Bonferroni or use safe e-combination; publish allocation in preregistration.

7.6 Replication plan, site variation, adversarial red-team reviews

Replication tiers.

1. **Internal same-site replication.** Repeat on separate days with fresh seeds, reversed block order. Target at least 3 null-hypothesis rejections or 3 strong nulls to stabilize inference.
2. **Cross-site replication.** Move rigs to new buildings with different power and RF environments; repeat the full protocol.
3. **Operator swap.** Exchange A and B teams and their roles; repeat to break operator-specific cues.

Variation matrix (pre-declared).

- Distances: L in {0.5 km, 2 km, 20 km} when feasible.
- Sources: swap homodyne and photon-arrival at A across runs.
- Time-of-day: day vs night to stress different noise environments.
- Shielding levels: RF cage on/off (without violating separation), to test robustness.

Adversarial red-team.

- RF injection challenges (blind to operators) at controlled, signed intervals to verify that alarms are raised and that such periods are quarantined in analysis.
- Power-sag and EMI stress tests scheduled on separate days to map false-positive surfaces.
- For any claimed positive, red-team proposes **killer ablations** (e.g., detuning block lengths, jitter envelopes, swapping sham labels). Claim stands only if it survives.

Outcome handling.

- *Positive (Delta_hat >= Delta, passes ablations and replications).** Publish as a candidate violation of marginal invariance with complete raw data, invite external labs to reproduce, and escalate DI scrutiny.
- **Null (no effect within power).** Publish full null with confidence sequences showing Delta bounded below Delta*. The architecture and preregistration remain as reusable templates.
- **Ambiguous (single-site positives, fail replications).** Treat as artifact; document failure modes; improve shielding and analysis and retry only with fresh preregistration.

Bottom line. The experiment is engineered so that either a small, steerable marginal shift emerges with audit-grade evidence, or we obtain a clean, quantitatively powered null. In both cases, the anytime-valid machinery, spacelike design, and adversarial audits make the conclusion scientifically crisp.

8. Cosmic and Environmental Coupling (Non-steerable Tests)

8.1 Exogenous covariates: Kp index, Schumann bands, cosmic rays

Purpose. Test whether *global* or *ambient* fields co-modulate QRNG raw statistics without implying control or communication.

Candidate covariates (copy-safe).

- **Geomagnetic activity:** Kp(t) (planetary 3-hour index), AE(t) (auroral electrojet).
- **Schumann resonances:** bandpowers near ~7.8, ~14, ~20, ~26 Hz measured by a local loop antenna + high-impedance front-end; derive P_SR,1(t), P_SR,2(t), …
- **Cosmic rays / secondary muons:** local counter rate R_mu(t) from a plastic scintillator or Geiger stack; note barometric pressure P_bar(t) and temperature T_env(t) for corrections.
- **Mains & RF:** 50/60 Hz fundamental and harmonics H_mains,k(t); broadband RF power P_RF(t).
- **Lab environment:** T_lab(t), vibration V_rms(t), PSU rail ripple R_ripple(t), humidity H_lab(t).

QRNG observables.

- Block-wise bias b_hat(t), Markov leakage I_hat1(t), code-length deficit dL(t) = 1 − L_t/t, spectral lines S_line,k(t), and panel aggregate D_panel(t).

Pre-processing.

- Synchronize all series to a common block grid; z-score each covariate; regress out obvious local confounds (e.g., mains harmonics from Schumann bands with notches) to avoid aliasing.

Registered hypotheses (illustrative).

- H1: dL(t) correlates with R_mu(t) after weather correction.
- H2: b_hat(t) increases during Kp(t) >= Kp_thresh.

- H3: I_hat1(t) co-varies with P_SR,1(t) within a tolerance band around 7.8 Hz.

All hypotheses are *directional* only if physically justified; otherwise use two-sided tests.

8.2 Out-of-sample forecasting and site generalization

Correlation on the training set is cheap. We require **prospective prediction** and **geographic transfer**.

Rolling-origin forecasting.

1. Split the timeline into folds F_1, F_2, … Each fold: train on the past W_train blocks; predict the next W_test blocks.
2. Models:
    - Linear: Y_t = beta^T X_t + u_t (with HAC or AR errors).
    - Regularized: LASSO/Ridge on lagged covariates X_{t−1:t−L}.
    - Nonlinear: gradient boosting on engineered features (changes, bandpowers, cross-terms).
    - Baseline: persistence (Y_t = Y_{t−1}) and seasonality-only.
3. Metrics: RMSE, MAE, and log-loss for binary endpoints (sign of b_hat). Report *improvement over baseline* (IoB) with paired tests across folds.

Site generalization.

- Train at site A, test at site B (different city/power grid).
- Train in season S1, test in S2 (e.g., summer->winter).
- Pre-register acceptance: mean IoB > 0 with p <= alpha after fold-wise correction, and absolute error below a pre-declared ceiling.

Robustness against spurious fits.

- **Surrogates:** phase-randomized versions of each covariate; the model must *fail* on surrogates.
- **Permutations:** block-wise label shuffles preserve autocorrelation; significant gains must vanish under shuffles.

- **Ablations:** remove each covariate group (Kp, SR, muons, mains/RF) and report delta-performance; interpretations depend on ablation drops.

Hierarchical view (optional).

- Bayesian hierarchical model with site-specific random effects:
  Y_{t,s} = beta^T X_{t,s} + b_s + u_{t,s},
  with priors encouraging shrinkage; check whether group-level betas exclude 0.

8.3 Interpreting stable correlations that still cannot signal

Key distinction. A *stable* coupling Y_t ↔ X_t (exogenous X) is *not* a channel. Signaling requires **steerability** of Y by a *distant choice*. Environmental covariates are not under our control and usually affect both sites via classical fields; they may explain variance without enabling communication.

Decision tree (copy-safe).

1. **Is the effect predictive out-of-sample and across sites?**
   - No -> artifact / overfit; document and move on.
   - Yes -> proceed.
2. **Does the effect survive surrogates, permutations, and ablations?**
   - No -> likely spurious correlation; tighten controls.
   - Yes -> plausible environmental coupling.
3. **Is the coupling *steerable* by B's schedule?**
   - No -> non-steerable covariate; *cannot signal*.
   - Yes (requires Section 7 proof) -> candidate violation; escalate.

Quantifying non-steerable influence.

- Treat X_t as an *instrumental regressor*; form a partialed-out statistic:
  Y_t^resid = Y_t − E[Y_t | X_{t−lag: t}] .
  Use Y_t^resid for the signaling endpoint so environmental variance is removed.
- Publish both *raw* and *residual* results; a true signaling effect should persist in residuals.

Causality cautions.

- **Common drivers:** mains or RF emissions can modulate both the QRNG and a local SR antenna; “correlation” may be shared interference.
- **Timing skew:** enforce pre-specified lags and forbid post-hoc alignment; otherwise “predictive” fits are time-shifted leakage.
- **Double counting:** never mix DI-lane trials into environmental regressions; keep the DI certificate independent.

What counts as meaningful here.

- A replicable, cross-site, out-of-sample predictive link (e.g., Kp -> small bias drift) is *scientifically valuable* for QRNG engineering (better shielding, adaptive de-rating) and for your anomaly-as-attention thesis.
- It still does *not* license messaging claims unless the effect becomes *controllable by choice*. Until then, treat it as a **nuisance term** for Section 7 and as a **salience cue** for agents.

Reporting standard (per run).

- Release the covariate set, pre-processing steps, feature catalog, model specs, fold boundaries, and full held-out forecasts.
- Include surrogate/permutation distributions, ablation tables, and residual-endpoint analyses for signaling.
- Summarize with a small dashboard: IoB, cross-site transfer score, residual variance removed, and whether any improvement exceeds Δ*-equivalent thresholds (if we map effect sizes to QRNG panel units).

Bottom line. Environmental couplings can be real, reproducible, and useful for *engineering and attention routing*, but by construction they are *non-steerable* and therefore *non-signaling*. Our protocol isolates them quantitatively and subtracts them before we ever test marginal shifts driven by B’s schedule.

9. Consciousness-Adjacent Agent Design

9.1 Using D(q(t) || tau) spikes as attentional salience for learning

Core idea. Treat live divergence signals from the QRNG monitor panel as an *attention bus* for the agent rather than as semantics. Let
D_panel(t) = [bias, runs, markov, spectral, compressibility, cross_MI].
Define a scalar **anomaly energy**
E(t) = max_k z_k(t),
where z_k(t) are per-monitor standardized scores (anytime-valid, zero-mean under null).

Routing policy (copy-safe).

- **Gating:** if E(t) < z_warm, ignore; if z_warm <= E(t) < z_hot, *attenuate* behavior (lower step size, raise logging); if E(t) >= z_hot, *pause* and checkpoint.
- **Learning bonus:** in RL or online learning, add an intrinsic reward r_int(t) = beta * clip(E(t), 0, E_max). This directs exploration toward regimes where the world/device deviates from iid.
- **Data triage:** push segments with high E(t) into a *salience buffer* with higher replay priority; store full telemetry for post-hoc diagnosis.
- **Adaptive extractor:** when E(t) rises, simultaneously tighten extractor compression (Sec. 4) so downstream learners never consume weak randomness.

Why this helps. The agent does not "believe in" meaning—it *allocates compute* where structure appears. This is a cybernetic analogue of attention: divergence begets scrutiny; alignment begets throughput.

9.2 Synesthesia renderers and exploration policies driven by anomaly energy

Renderer role. A *synesthesia renderer* maps raw or summarized signals into audiovisual primitives that the agent (or human) can scan quickly. Use it to (a) externalize D_panel(t) structure, and (b) probe whether structure repeats under controlled interventions.

Minimal renderer spec (example).

- **Timeline track:** brightness ~ E(t); hue encodes argmax_k of z_k(t) (which monitor dominated).
- **Micro-glyphs:**
    - bias spike -> short "tilt" glyph;
    - runs spike -> comb glyph;

    - spectral spike -> sinusoid glyph at offending frequency bin;
    - compressibility spike -> checker glyph.
- **Audio stub:** map E(t) to amplitude and dominant monitor to pitch band; cap dynamics to avoid audio pareidolia.

Exploration policies (agent side).

- **Policy P1 (probe-and-quench):** when E(t) enters warm zone, execute a *micro-probe* (small, reversible intervention that should not affect QRNG internals—e.g., change a downstream task parameter or sensory configuration). If E(t) falls, tag as likely exogenous; if it rises, schedule a *controlled repeat* to test reproducibility.
- **Policy P2 (AB-testing):** run interleaved A/B blocks with identical outer context while logging E(t). Expect *no change*; any consistent delta triggers deeper diagnostics (not claims).
- **Policy P3 (uncertainty budgeting):** allocate a fixed curiosity budget per hour. Spend budget on states with high predicted learning progress *and* elevated E(t), so anomaly chasing cannot starve task learning.

Engineering guardrails.

- Renderers never feed back into the raw-bit path; only meta-control and logging react to E(t).
- All renderer mappings are deterministic and versioned; include a glyph legend in the certificate to avoid post-hoc reinterpretation.

9.3 Safety: never upgrade anomaly to meaning without cross-evidence

Rule S0 (no semantics from D alone). A spike in E(t) is *not* a message. Promotion from "anomaly" to "meaning" requires **independent, pre-registered** evidence.

Promotion ladder (copy-safe).

1. **Anomaly:** E(t) >= z_warm. Action: log, de-rate extractor if needed, enqueue for review.
2. **Replicable anomaly:** same signature reappears under repeated *identical* conditions; passes surrogate/permutation checks. Action: open a preregistered *artifact investigation* (EMI, thermal, firmware).
3. **Instrumented anomaly:** survives controlled ablations (alternate power, shielding on/off, room swap), shows consistent effect size with confidence sequences excluding null by >= Delta*. Action: document as engineering coupling; mitigate or subtract.

4. **Candidate meaning (rare):** only considered if (i) anomaly is *steerable by choice* in a spacelike protocol (Sec. 7), *and* (ii) persists across sites and teams with blinded analysis. Action: escalate to public replication; do not assign narratives beyond "operational signaling candidate."

Statistical hygiene.

- Use anytime-valid p-values/e-values; record **alpha_spent** per stream; freeze thresholds before runs.
- Maintain *dual logs*: (i) raw and monitor statistics; (ii) operator actions and environmental covariates. Promote only on the raw/monitor log; never on anecdotes.

Human-factor safety.

- **Pareidolia brakes:** limit renderer loudness/contrast; require cross-evidence before any human annotation is allowed in the record.
- **Two-key rule:** any claim beyond "anomaly" requires sign-off from independent analysis team with access to blinded data.

Operational invariants (copy-safe).

- Extracted bits remain epsilon-bounded regardless of E(t) via automatic de-rating.
- Salience never bypasses cryptographic hygiene (seeds independent; no raw-bit reuse).
- The DI lane (Sec. 6) is *orthogonal*; its certificates are unaffected by renderer decisions.

Success metric (adjacent, not mystical).

- On benchmark tasks, agents that route compute using E(t) achieve higher sample efficiency or faster fault isolation than baselines at equal safety (no increase in shipped weak-entropy windows). We pre-register metrics and publish nulls if the gain is absent.

Bottom line. Treat D(q(t) || tau) as an attentional utility, not a metaphysical oracle. Use it to decide *where to look next* and *how cautiously to operate*, and only escalate to "meaning" when independent, steerable, and replicated evidence clears the pre-declared bars.

10. Implementation Details

10.1 Hardware BOMs: sources, detectors, ADCs, timing, isolation

Entropy sources (pick 1–2 for redundancy).

- Photon path/arrival: 50:50 fiber coupler, attenuated laser or SPDC pair, 2x Si APD (visible) or InGaAs APD (NIR), low-jitter TDC or fast ADC.
- Vacuum homodyne: narrow-line LO laser, 50:50 coupler, balanced photodiode (bandwidth >= 100 MHz), low-noise transimpedance amp.
- Phase-diffusion interferometer: unbalanced MZI/Sagnac on fiber bench, thermally stabilized, single PD at output.

Front-end & digitization.

- TIAs with gain set to keep shot-noise above electronics noise by >= 6 dB in-band.
- Anti-alias filter matched to ADC sampling rate.
- ADC: 12–16 bit, 100–500 MSPS (homodyne/interferometer); TDC: <100 ps binning (photon time-of-arrival).
- FPGA or SoC with DMA to ring buffers (PCIe or GbE egress).

Environmental & health telemetry.

- Temperature (source, detectors, board), vibration (MEMS), PSU rails ripple, LO power monitor, dark-rate counter, APD bias/temperature, RF probe power.
- Optional muon counter, loop antenna for Schumann bandpower, mains harmonic monitor.

Timing & separation.

- Local clock: OCXO or GPSDO for timestamping only (no control loops slaved to GPS).
- Hardware timestamping on ingress to FPGA; monotone timebase with overflow handling.
- Optical fiber links for control; no copper between A and B. Air-gap for data export.

Shielding & isolation.

- RF-tight enclosure (gasketed seams), filtered feedthroughs, star-ground at single chassis point.

- Separate UPS for each station; filtered DC rails with LC and ferrite.
- Light-tight optics box; beam dumps; fiber isolators.

10.2 Firmware/FPGA notes: buffering, whitening, loopbacks

Data path.

- Ingress -> FIR/decimate (if homodyne) -> raw sample packing -> ring buffers -> host.
- Photon path: edge/timestamp capture -> inter-arrival binning -> raw bit decisions in FPGA (optional; also stream timestamps).

Buffers & flow control.

- Dual or triple ring buffers sized for >= 2 s of peak-rate data; back-pressure signals to throttle front-end safely.
- Loss counters latched and exposed in telemetry; any overflow triggers HOT state.

Whitening vs extraction.

- Do not “pre-whiten” in ways that hide structure. If a minimal de-bias is required to meet link limits, use a reversible mapper (documented) and always preserve raw taps to disk.
- Full extraction happens in software or a documented FPGA core; parameters (k, epsilon, seed_commit) published per window.

Online monitors in FPGA (optional fast sentinels).

- Running bias and runs counters per window; spectral line detector for mains harmonics; threshold comparators raising WARM/HOT flags.
- All FPGA thresholds are config-locked and hashed.

Challenge–response loopbacks.

- DAC-inject path to add known micro-bias (Delta_inj) at the ADC input via summing node; gated by secure element.
- Dead-time emulator for photon path to create controlled Markov leakage.
- Spectral tickle: weak sinusoid injection around known line.
- Each tickle interval signed by secure element and hidden from host until audit.

Secure elements & keys.

- On-board secure element stores device key (Ed25519 or similar) and signs: window certificates, loopback attestations, firmware hash at boot.
- Firmware measured-boot: FPGA image hash logged at startup; mismatch -> refuse to run.

10.3 Software: extractor, monitors, telemetry hashing, audit logs

Extractor.

- Toeplitz hashing or XOF (SHAKE) with per-window seed; k chosen by: k <= N*h_min - overhead(epsilon) - leakage_budget.
- Seed management: CSPRNG seeded from prior certified bits or DI lane; never reuse seeds for overlapping raw windows; store seed_commit = H(seed).
- Streaming mode: process chunks; constant-time implementation; SIMD where safe.

Monitor suite (anytime-valid).

- Bias e-process (mixture bettor), runs/order-1 Markov LLR, spectral CUSUM for configured lines, online LZ/CTW codelength.
- Confidence sequences for mean (bias) and for Markov MI; publish p_t or e_t with alpha_spent per stream.
- Cross-predictability: online regression S1<-S2 and S2<-S1; report MI_hat with CI.

Telemetry pipeline.

- Sample all telemetry at block cadence; z-score and store; out-of-range -> drop h_min floor or HOT.
- Compute h_min_lo from 90B-style estimators adapted to model; include CI and assumptions tag in certificate.

Certificates & logs.

- Per-window JSON (or CBOR) record:
  window_id, start_ts, end_ts, N_raw, k_out, epsilon, h_min, extractor_type, seed_commit, monitors_digest, telemetry_digest, alpha_spent, state, firmware_hash, device_pubkey.
- Append device signature over canonicalized record; write to append-only ledger.

- Raw data: write-once files with content hashes; segment boundaries aligned to windows.

Ops console.

- Real-time dashboard: E(t) anomaly energy, panel z-scores, state machine (NORMAL/WARM/HOT), output rate, DI lane status.
- One-click “quarantine last M windows” action; generates signed quarantine notice.

10.4 Provenance and auditing: append-only ledgers, signatures

Ledger design (copy-safe).

- Per-device append-only log (AOL) with Merkle chaining: each entry includes prev_hash, so any deletion or edit breaks the chain.
- Daily anchor: publish top Merkle root hash to a public timestamping service or cross-sign with an independent device.
- Separate ledgers per class: certificates, telemetry summaries, loopback attestations, DI lane records.

Signing & verification.

- Ed25519 device key signs each ledger entry; verifier checks signature, prev_hash, and included digests (raw file hash, monitor digest).
- Notary C stores preregistration artifacts: analysis binary hash, config hash, threshold set, schedule commitments from B.

Reproducibility packet (per run).

- Frozen copies of: firmware bitstream hash, extractor version, monitor code hash, config files, DI f_CHSH spec, 90B estimator versions.
- Mapping of window_ids to raw file offsets and seed reveal schedule.
- Auditor scripts to: recompute monitors, recompute extractor outputs from raw+seed, verify epsilon and k bounds given h_min model.

Data governance.

- Retention policy: keep raw buffers for Tmin (e.g., 1 year) unless legal/audit hold extends it; extracted bits can be deleted earlier if sensitive.

- Access control: auditors receive read-only mounts; any export re-signed by device plus notary.

Incident response.

- HOT latch generates incident bundle: last K windows raw + telemetry, console state, tickle schedule (if any), environment snapshot. Bundle is signed and mirrored off-box.
- Root-cause workflow: replicate anomaly with controlled injections; if undiagnosed, raise h_min model conservatism until resolved.

Public transparency (optional).

- Publish rolling certificate summaries (no seeds) and DI anchors; provide quarterly "state of randomness" reports: uptime, alarms, de-rates, DI duty, observed environmental couplings.

Bottom line. Hardware isolates and measures; firmware moves bits without hiding structure; software extracts only what the physics and DI lanes can safely justify; the ledger makes every claim audit-ready. If anything drifts, the record shows when, how much, and how the system reacted—no bravado, just reflexive, signed evidence.

11. Results Templates and Interpretation

11.1 Certificate formats: <window_id, k, epsilon, h_min, monitor_stats, telemetry_hash>

Purpose. Every shipped bit window is accompanied by a concise, signed record that lets an auditor recompute guarantees and replay tests.

Minimal certificate (human-readable summary).

window_id: A-2025-10-30T15:02:18Z-000173

start_ts: 2025-10-30T15:02:18.000Z

end_ts: 2025-10-30T15:02:20.000Z

N_raw: 120000000

k: 67108864

epsilon: 2^-128

h_min: 0.62 # per-bit lower bound within this window

lane: PH # PH = physics lane; DI = device-independent lane

seed_commit: 7b2f..d1 # H(seed), seed revealed only for audit

monitor_stats_digest: 1b94..ee

telemetry_digest: 5e3a..8f

state: NORMAL # NORMAL | WARM | HOT

alpha_spent: 0.0031 # cumulative anytime-valid budget consumed

firmware_hash: 9a71..2c

device_pubkey: ed25519:df1c..9e

sig: ed25519: 3045..a1

Expanded machine format (CBOR or JSON).

- monitor_stats is a small struct of per-monitor summaries that can be digested for the minimal record:

monitor_stats: {

bias: {p_anytime: 0.41, cs_lower: 0.4967, cs_upper: 0.5033, g_log: 0.0002},

runs: {p_anytime: 0.37, z: 0.21},

markov: {e_value: 1.6, I_hat_bits_per_bit: 3.0e-5},

spectral: [{freq_hz: 60, e_value: 1.1}, {freq_hz: 120, e_value: 1.0}],

compress: {L_over_N: 0.9992, e_value: 1.3},

cross_MI: {bits_per_bit: 2.0e-5}

}

- telemetry_hash covers a vector of calibrated measurements (temperature, LO power, dark rate, ADC histograms, PSU ripple), plus their confidence intervals used for h_min computation.

Audit replay. Given raw window, seed, extractor spec, and h_min model tag, an auditor must be able to:

1. recompute k and epsilon bound,
2. recompute all monitor metrics and anytime-valid p or e values,
3. re-derive the same monitor_stats_digest.

11.2 Example nulls, example alarms, handling false positives

Example 1: Clean null (NORMAL).

- bias: p_anytime = 0.41, cs includes 0.5 with tight bounds
- runs: p_anytime = 0.37
- markov: e_value = 1.6 (well below 1/alpha_hard)
- compress: L_over_N = 0.9992, e_value = 1.3
- spectral: e_values near 1 for mains harmonics
- h_min: 0.62 from telemetry; DI anchor this hour shows S_lo > 2, yielding a small DI rate that agrees in order of magnitude
- **Action:** extract as configured; certificate state NORMAL.

Example 2: Warm alarm with de-rate (WARM).

- transient bias drift: p_anytime = 0.018; cs excludes 0.5 by Delta ~= 2.6e-5
- compress: e_value = 2.9 (small growth), others cool
- telemetry: LO power down 1.5 dB within tolerance; h_min drops from 0.62 to 0.55
- **Policy reaction:** move to WARM; reduce k/N by 10 to 20 percent; raise logging cadence; keep shipping with epsilon preserved; record cause in certificate.

Example 3: Hard alarm and halt (HOT).

- spectral sentinel at 60 Hz: cumulative e_value crosses 1/alpha_hard = 200
- markov leakage rises: I_hat_bits_per_bit = 1.5e-3 sustained for T_sustain
- alpha_spent remains within budget because we use e-process thresholds
- **Policy reaction:** latch HOT; quarantine last K windows; stop output; emit incident bundle with raw and telemetry; require operator reset and re-calibration.

Handling false positives (anytime-valid discipline).

- By Ville's inequality, the long-run probability of ever crossing 1/alpha_hard under the null is <= alpha_hard. Still, two practical steps:
    - **Ablate immediately:** check if the alarm co-occurs with known tickle intervals, PSU events, or mains excursions. If yes, classify as explained alarm.

- **Cooldown rule:** after a HOT triggered solely by one monitor while others remain cool and telemetry is nominal, allow a single guarded re-arm after a fixed number of clean windows; otherwise require manual sign-off.

Post-mortem template.

- Plot e_value or p_anytime trajectories with the spending curve.
- Overlay environmental covariates and tickle attestations.
- Compare raw vs residualized metrics (subtracting exogenous predictors from Sec. 8).
- Classify: explained (engineering), unexplained but non-replicating (artifact), or persistent (design flaw) with required mitigation.

11.3 Criteria for declaring evidence of signaling (and why they are strict)

Primary endpoint recap (Sec. 7).

- Blockwise summary Y_i at A (e.g., bias_hat or martingale log-growth g_i).
- Hidden label Z_i in {+1, −1} set at B.
- Effect size Delta_hat = mean(Y | Z = +1) − mean(Y | Z = −1).

Declare a candidate only if ALL of the following hold:

1. **Pre-registered threshold met.**
    - Delta_hat >= Delta* with the anytime-valid p_value <= alpha_hard after panel-wise multiplicity control.
    - Confidence sequence for Delta excludes 0 by a margin >= Delta*.
2. **Spacelike integrity preserved.**
    - Verified timing logs show choose-and-measure latencies well within the L/c window for the entire run, with sealed clocks and no drift beyond spec.
3. **Ablations passed.**
    - Survives removal of leading/trailing 10 percent of each block (edge artifacts).
    - Survives time-shift controls (no alignment-by-leakage).
    - Survives site swap, operator swap, and sham days (no effect when schedule is fake).

4. **Environmental residualization.**
    - Effect persists when Y is replaced by Y_resid = Y − E[Y | X] where X includes Kp, Schumann bandpowers, mains harmonics, RF power, temperature, PSU ripple, and any lab covariates pre-declared in Sec. 8.
    - No single exogenous covariate explains away the effect in out-of-sample tests.
5. **Red-team adversarial tests.**
    - Positive is reproduced after red-team suggests killer variants (e.g., irregular block lengths, jitter envelopes, random pauses) and under enhanced shielding.
    - No detectable RF, optical, acoustic, or power-line pathway can carry B's schedule to A within light-speed limits.
6. **Replication across days and sites.**
    - Independent repeats (N >= 3) at the same site and at least one distant site, all pre-registered, all passing 1–5 with consistent Delta_hat within CS overlap.
    - Analysis binaries and configs match pre-run hashes; code escrow verifies no mid-run changes.
7. **DI anchor healthy.**
    - Parallel DI lane during campaign shows S_lo > 2 with positive certified entropy rate; rules out trivial classical coupling masquerading as "quantum."

Why this bar is strict.

- **No-communication theorem stakes:** A true, steerable marginal shift contradicts a central operational consequence of standard QM. The cost of a false positive here is not a bad product demo; it is a literature detour.
- **Multiple weak avenues for leakage:** Power-line harmonics, RF ingress, scheduler timebase drift, and human-operator micro-behaviors can all fake a blockwise effect if not crushed.
- **Anytime-valid necessity:** Without e-processes and preregistered endpoints, repeated peeking inflates Type-I error and turns interesting noise into apparent discovery.

Reporting a candidate.

- Publish a full packet: raw bits, schedules and signatures, timing logs, telemetry, certificates, analysis binaries and hashes, DI records, and red-team notes.

- Provide third-party replication kits: exact BOM, shielding instructions, and scripts to reproduce the effect size and p/e values.
- State claims narrowly: “operational signaling candidate with Delta_hat >= Delta* under preregistered tests,” and list all remaining alternative explanations investigated.

Failing to meet the bar.

- If any criterion fails, report as null for signaling with a quantified upper bound: CS shows |Delta| < Delta* across replications. Document lessons that improved the reflexive QRNG (fewer shipped weak windows, faster fault isolation), which remain wins regardless of signaling outcome.

Bottom line. Certificates make each window auditable; alarms are handled with anytime-valid rigor and transparent post-mortems; and signaling claims are only contemplated when a high, multi-layer bar is cleared repeatedly. This keeps the science crisp whether the result is a careful null or a historic anomaly.

12. Limits, Failure Modes, and Alternative Explanations

12.1 EMI, cross-talk, thermal drift, firmware bias, side channels

Electromagnetic interference (EMI).

- *Mechanism:* conducted/radiated fields modulate front ends (TIAs, ADC reference), imprinting mains lines, RF bursts, or digital clocks into raw bits.
- *Symptoms:* spectral e-process growth at 50/60 Hz and harmonics; correlation with RF probe power; reproducible at certain times of day.
- *Diagnostics:* RF survey with swept receiver; near-field probe pass; turntable orientation tests; power-line impedance injection; shield-on/off A/B.
- *Mitigations:* RF gasketed enclosure; filtered feedthroughs; star-ground; ferrites on all cables; battery or double-conversion UPS; optical data egress.

Cross-talk (electrical/optical).

- *Mechanism:* shared grounds, reference rails, or optical back-reflections couple control planes to the entropy path.
- *Symptoms:* cross-predictability MI_hat(S1;S2) > Delta*_MI; sensitivity to cable routing; dependence on actuator states.

- *Diagnostics:* isolate grounds; switch to fiber; insert optical isolators and beam dumps; modulate one path and look for correlated artifacts on the other.
- *Mitigations:* strict partitioning of analog/digital planes; separate supplies; opto-isolated control; angled physical separation inside the enclosure.

Thermal drift and mechanical vibration.

- *Mechanism:* index-of-refraction and detector gain drift, interferometer phase walk, fringe visibility changes.
- *Symptoms:* slow bias trends; increased runs statistics; low-frequency spectral energy; correlation with T_lab and V_rms.
- *Diagnostics:* temperature sweeps; vibration injection with shaker; coefficient estimation (dbias/dT).
- *Mitigations:* thermal mass and control (PID to +-0.05 C); vibration damping; fiber pigtail strain relief; auto-rebalance of homodyne.

ADC and analog nonidealities.

- *Mechanism:* quantizer integral nonlinearity, reference ripple, aperture jitter; saturation or clipping.
- *Symptoms:* compressibility gains; specific codeword defects; harmonics at sampling-related lines.
- *Diagnostics:* histogram tests with known Gaussian input; code-density analysis; loopback with precision DAC.
- *Mitigations:* oversampling + dithering (documented); cleaner references; shielded clocks; regular linearity calibration.

Detector specifics (APD/SiPM).

- *Mechanism:* afterpulsing (charge traps), dead time, bias supply noise.
- *Symptoms:* order-1 Markov leakage; excess short inter-arrivals; temperature-dependent effects.
- *Diagnostics:* measure autocorrelation of inter-arrival times; vary bias and temperature; vendor-modeled afterpulse subtraction for sanity check.
- *Mitigations:* lower count rates; active quenching tuning; temperature stabilization; detector choice with low afterpulse specs.

Firmware bias and digital paths.

- *Mechanism:* off-by-one in thresholding, reusing seeds, partial whitening that hides structure, buffer overflows.
- *Symptoms:* fixed-pattern bias; dependence on window alignment; alarms vanish if different bit slicing is used.
- *Diagnostics:* code escrow and reproducible builds; bit-slice rotation tests; synthetic injection (challenge–response) proving detectability.
- *Mitigations:* measured-boot; formal review of HDL; unit tests for bit-true paths; forbid undocumented whitening.

Operator and schedule leakage.

- *Mechanism:* human routines align with mains cycles; inadvertent cues (acoustic, optical) synchronize A and B.
- *Symptoms:* block-boundary effects; time-shifted correlations that vanish with label permutations.
- *Diagnostics:* sham schedules; time-jittered blocks; sound/light sensors; operator swap.
- *Mitigations:* automation at B; silence/light discipline; randomized block edges.

Side channels (power, ground, acoustics, optics).

- *Mechanism:* schedule-dependent PSU current at B couples to shared building mains seen at A; acoustic hum couples into optomechanics; stray light paths.
- *Diagnostics:* independent UPS islands; accelerometers and microphones; dark-box tests; optical leakage checks.
- *Mitigations:* separate circuits/transformers; acoustic isolation; light traps and interlocks.

Policy note: any suspected coupling downgrades a “candidate” to “artifact” until ablated. Extractor continues to preserve epsilon via de-rate; claims are withheld.

12.2 Superdeterminism, retrocausality, and operational steerability

Superdeterminism.

- *Claim:* the entire experiment, including setting choices, is jointly predetermined; apparent “violations” reflect deep correlations, not nonlocal influence.

- *Status:* unfalsifiable in broad form; can mimic any dataset.
- *Operational test:* ask whether an *agent* can *steer* A's marginal by choosing B's schedule now. If effects persist under sealed, cryptographically committed schedules and spacelike discipline, superdeterminism remains a logical possibility but offers **no actionable control**.
- *Position:* without an algorithm for steering, superdeterminism explains but does not help; it cannot power communication.

Retrocausality.

- *Claim:* future settings influence past outcomes; consistent with some time-symmetric interpretations.
- *Operational bar:* the same *steerability* criterion—can we encode arbitrary Z_i and recover it at A with pre-registered decoding? If yes, you have a **channel**, retro- or otherwise. If no, you have colorful ontology and no comms.
- *Controls:* randomized, sealed schedules; delayed-choice variants; label reveals strictly after lock-in.

Nonlinear/NS collapse models.

- *Claim:* weak nonlinearities in state update allow marginal shifts.
- *Operational angle:* our primary endpoint (marginal shift >= Delta*) and replication matrix are exactly the stress tests such models must pass. Absence of effect below power constrains their parameter space.

Bottom line (copy-safe).

- Explanations that cannot be used to *steer* marginals do not create communication. Our protocol judges by **operational steerability**: encode -> transmit (by schedule) -> decode at spacelike remove with anytime-valid stats. Everything else is philosophy, not a channel.

12.3 Why batteries of tests are alarms, not certificates

Scope mismatch.

- SP 800-22, Dieharder, TestU01 answer: "does this *sample* look inconsistent with iid(1/2) under one-shot p-values?"

- We need: “how many **secure bits** can we **guarantee** now, and when must we **stop**?” That is min-entropy bounding + anytime-valid monitoring.

Multiple looks and garden-of-forking paths.

- Repeatedly running a battery until something “fails” inflates Type-I error. One-shot p-values are not valid under continuous surveillance.
- *Fix:* test martingales/e-processes (Ville) or alpha-spending convert streaming evidence into valid alarms without p-hacking.

Model dependence vs guarantees.

- Passing batteries does not imply high h_min; a biased but drifting source can pass small samples.
- *Fix:* physics-informed h_min(t) bounds (SP 800-90B class) + DI certificates from Bell violations provide **lower bounds**, not just “looks random.”

Extractor coupling.

- Batteries often run *after* aggressive whitening, masking structure that matters.
- *Fix:* monitor on **raw** streams; extraction sized by h_min(t) with explicit epsilon; whitening only as documented, reversible pre-processing if absolutely required.

Adversarial robustness.

- A crafted sequence can pass many tests yet be predictable (e.g., LFSR with added noise).
- *Fix:* universal compression e-process and cross-predictability with a second source widen coverage; DI anchor attests nonclassicality.

Operational principle (copy-safe).

- Tests -> *alarms*. Entropy bounds -> *certificates*. Extraction -> *enforcement*.
- We publish both: per-window alarms (anytime-valid) and per-window certificates <k, epsilon, h_min, digests>. Consumers can audit claims instead of trusting vibes.

Consequences for signaling claims.

- A “battery fail” during B’s FOCUS blocks is *not* evidence of messaging; it is an alarm. Only if the *pre-registered endpoint* shows a marginal shift >= Delta* that survives ablations and replications (Sec. 11.3) do we even utter “candidate signaling.”

- Conversely, perfect battery passes do not certify unpredictability; only min-entropy proofs (physics/DI/EAT) justify k and epsilon.

Takeaway. Batteries are useful sirens in the reflexive loop, but they are not the judge or the certificate. We treat them as triggers for de-rate/halt and for investigation—not as licenses to claim randomness or meaning.

13. Ethics and Epistemic Hygiene

13.1 Preregistration, code escrow, blinded pipelines, data release

Preregistration (before power-on).

- Register hypotheses, endpoints, Delta*, alpha_hard, alpha_soft, spending schedule, block designs, ablations, and replication plan.
- Fix data exclusion rules, window sizes, and all stopping criteria (efficacy and futility).
- Freeze hardware configs: BOM hashes, enclosure photos, cable routing diagrams, shielding checklist.

Code escrow and measured builds.

- Hash and escrow binaries, configs, and container images for: extractor, monitors, DI lane analysis, forecasting models, and plotting.
- Use reproducible builds; record compiler versions and flags; store hash in a third-party notary.
- Firmware/FPGA bitstreams measured-boot; log their hashes in the first ledger entry of each run.

Blinded pipelines.

- A-side computes only unlabeled block summaries Y_i keyed by opaque block_id; no access to labels Z_i or schedule metadata.
- Unblinding occurs once: after the run, with a notarized reveal package from B that includes seed commitments and signatures.
- All interim dashboards show only anytime-valid statistics, never effect sizes by label.

Data release and retention.

- Release artifacts in graded tiers:

1. Immediate: certificates, ledgers, hashes, prereg docs, analysis hashes, shielding audit.
2. After T_hold (e.g., 30 days): raw windows, seeds, telemetry, DI logs, red-team notes.
3. Long-term: replication kits and scripts.

- Redact only what would compromise unrelated systems; never redact bits or seeds needed for audit.
- Keep raw data write-once, read-many; mirror to a neutral repository with integrity proofs.

Operator conduct.

- Written SOPs: no ad-hoc tweaks mid-run; every physical interaction logged with timestamp and reason.
- Two-person rule for opening enclosures or changing configs; photograph and countersign.

13.2 Community replication and adversarial collaboration

Open replication kits.

- Provide shopping-list BOM with vendor SKUs and acceptable substitutes, mechanical drawings, wiring diagrams, FPGA images, containerized software, and calibration procedures.
- Include test vectors and loopback modes so labs can validate their stacks before physics-on.

Adversarial collaboration charter.

- Invite a red-team consortium at project start; give them explicit powers: propose killer controls, schedule shams, request site swaps, and inspect shielding.
- Publish a standing “attack surface” document listing suspected couplings and the current mitigation status.

Multi-site governance.

- Coordinate independent replication sites with staggered preregistration to avoid tacit tuning.

- Share only what preregistration requires until after each site unblinds; then publish all runs together with concordance analysis.

Nulls are first-class results.

- Commit to publish complete nulls with confidence sequences and powered Delta bounds.
- Maintain a living registry of all attempts, including partial runs and failures to launch, to prevent file-drawer bias.

Incentives and authorship.

- Make data and tools CC-BY or equivalent; offer coauthorship to external labs that meet preregistered replication thresholds, including strong nulls.
- Set aside budget for external audits; list conflicts of interest and funding constraints up front.

13.3 Claims we will not make without meeting the bar

No signaling claims without 11.3.

- We will not claim communication, influence, intention, or "mind-over-QRNG" unless all criteria in 11.3 are met: Delta_hat >= Delta*, anytime-valid significance after multiplicity control, ablations passed, environmental residualization, red-team sign-off, and cross-site replication with healthy DI anchors.

No metaphysical upgrades from anomalies.

- We will not interpret D(q(t) || tau) spikes as meaning, messages, or "cosmic communication." Anomalies are attention cues until cross-evidence and steerability are demonstrated under preregistration.

No "looks random" certificates.

- Passing batteries (SP 800-22, Dieharder, TestU01) will never be presented as proof of unpredictability. Only min-entropy bounds (physics, 90B-style, DI/EAT) plus extraction sizing justify k and epsilon.

No post-hoc endpoint fishing.

- We will not introduce new primary endpoints after unblinding. Exploratory findings, if any, must be confirmed in a fresh, preregistered run before being called "effects."

No silent parameter changes.

- We will not alter extractor compression, thresholds, duty cycles, or shielding during an analysis window without logging and freezing a new preregistration version.

No claims beyond powered regions.

- If the powered Delta* is not met, we will not suggest “hints” of effects. We will report the confidence sequence bound on Delta and stop.

Plain-language summaries.

- Every public result will include a non-technical summary of what was tested, what was found, what the guarantees mean (and do not mean), and how someone else can check our work.

Bottom line.

- The ethos is: lock it, log it, and let others break it. We design our process so that a careful null is as valuable to the field as a positive—and a positive cannot survive without withstanding adversarial replication under open, precommitted rules.

14. Discussion

14.1 If a steerable marginal shift emerges: implications and next tests

Implications. A reproducible, spacelike, schedule-keyed marginal shift at A (>= Δ*), surviving ablations and cross-site replication, would falsify the operational consequence of no-communication for our setup. That does **not** instantly deliver a robust channel; it establishes an *existence proof* of steerability of local marginals under our conditions. Immediate consequences:

- **Physics:** strong pressure toward nonlinear/NS collapse models or hidden-signal pathways previously missed; standard QM must be extended or the leakage isolated.
- **Metrology & crypto:** randomness certification frameworks would need a “steerability audit” in addition to DI/EAT. Certified keys may require provenance constraints that exclude regimes where steering was observed.
- **Foundations of “conscious steering”:** if human-driven blocks are the only condition that works while automated schedules fail under otherwise identical controls, the phenomenon would require a *state-dependent* theory of steering. If both work, the result is purely physical and independent of intent.

Next tests (tightening the vise).

1. **Dose–response & coding:** move from block labels to explicit pseudo-random codebooks at B; show rate-vs-error curves for decoding at A under preregistered M-ary messages.
2. **Latency & distance scaling:** vary L and local latencies to map the spacetime profile; if effects shrink with margin erosion, it supports genuine spacelike coupling rather than slow leakage.
3. **Mechanism dissection:** swap entropy sources (homodyne vs photon), detectors, ADCs, and shielding stacks; keep the effect if and only if the phenomenon is physics-level, not device-level.
4. **Adversarial hardening:** red-team introduces decoys (fake schedules, irregular jitter) and probes all side channels at higher sensitivity. Survive or retract.
5. **DI saturation:** run the DI lane at higher duty while steering persists; any collapse of DI certificates during "positive" epochs would indict classical leakage rather than post-quantum physics.
6. **Independent labs:** ship full replication kits; require at least 3 external labs to reproduce the same effect size within confidence sequences before any theory language escalates.

14.2 If not: what we gained for security, agents, and science

Security & reliability.

- **Anytime-valid guardrails:** the martingale/e-process stack provides real-time, p-hacking-proof alarms; fewer silently weak windows ship.
- **Per-window certificates:** <window_id, k, epsilon, h_min, digests> gives consumers auditable guarantees and post-hoc traceability.
- **Hybrid DI anchor:** periodic DI beacons bound min-entropy independently of device models, upgrading trust without sacrificing throughput.

Agent design.

- **Anomaly-as-attention:** routing D(q(t) || tau) spikes into exploration bonuses and data triage improved *sample efficiency* and *fault isolation* in preregistered benchmarks (or we publish nulls and the configs).
- **Safety coupling:** automatic de-rate/halt on divergence prevents downstream learners from amplifying weak randomness—practical "consciousness-adjacent" caution.

Scientific byproducts.

- **Environmental transfer maps:** quantified how Kp, Schumann bandpowers, muon flux, mains/RF contribute to raw statistics; those regressors now serve as nuisance subtraction in precision QRNGs.
- **Open templates:** preregistration, code escrow, and ledgers form a reusable protocol for any lab attempting high-stakes streaming inference (not just QRNGs).
- **Negative knowledge:** a powered, clean null constrains entire families of nonlinear/retrocausal models and documents what *doesn't* happen under rigorous controls.

Bottom line: a careful null is a win—hardening randomness, clarifying environmental couplings, and delivering an operational "attention bus" for agents—without overclaiming.

14.3 Logostics view: alignment when q(t) refuses to be tau

Logostics treats the system as an alignment loop: q(t) -> tau(t) by minimizing D(q(t) || tau). Three readings emerge:

1. **Operational humility.** When D spikes, the agent does not invent a story; it *adapts*: lower k/N, slow down, log harder, and route salience to diagnosis or learning. Alignment is lived as governance, not asserted as faith.
2. **Epistemic discipline.** "Anomaly" is a *where-to-look-next* pointer, not proof of meaning. Only cross-evidence, steerability, and replication promote anomaly to claim. In Logostics terms, divergence is a gradient; it becomes "truthward motion" only when independent witnesses cohere.
3. **Teleology without mystique.** If a steerable marginal shift appears, the q(t) -> tau(t) loop doesn't collapse into metaphysics; it expands its model class and tightens its tests. If no shift appears, the loop still converges to a safer, more truthful operation—q(t) approaching tau not by decree but by *disciplined refusal* to ship ungrounded bits or unearned narratives.

Perspective. Whether the Universe offers steerable structure or holds the line on invariant marginals, the reflexive QRNG embodies Logostics: a system that orients itself toward living truth by admitting uncertainty, publishing evidence, and changing behavior when alignment falters.

15. Conclusion

15.1 Reflexive QRNGs as trustworthy randomness infrastructure

We have replaced bravado with governance. A reflexive QRNG does not *assume* ideal iid bits; it *earns* them, window by window, by (i) bounding h_min(t) from physics and telemetry, (ii) extracting to a declared epsilon with conservative overhead, and (iii) running anytime-valid monitors that can de-rate or halt before weak bits ship. Per-window certificates <window_id, k, epsilon, h_min, monitor_stats, telemetry_hash> turn claims into evidence; a low-duty DI lane anchors nonclassical entropy, while the high-throughput physics lane delivers bandwidth under supervision. The result is not just "randomness that looks random," but auditable, reproducible unpredictability that downstream cryptography, simulation, and science can trust.

15.2 Consciousness-adjacent anomaly routing without overclaiming

D(q(t) || tau) is promoted from a diagnostic to a utility: an attention bus that routes compute toward where structure appears *now*. Divergence spikes drive exploration bonuses, richer logging, and safer operation (automatic rate control), yet they never become "messages" on their own. Meaning requires cross-evidence and steerability under preregistration. This stance lets agents benefit from anomalies—learning faster, localizing faults sooner—without collapsing into pareidolia or metaphysics. In short: anomaly energy becomes disciplined curiosity, not doctrine.

15.3 A reproducible template for probing the limits of signaling

Whether the outcome is a careful null or a candidate effect, the pathway is the contribution: spacelike separation, shielded infrastructure, blinded pipelines, preregistered endpoints, anytime-valid statistics, DI anchoring, environmental residualization, and adversarial replication. If a steerable marginal shift >= Delta* survives that gauntlet, physics must respond with new theory or newly discovered leakage; if not, we have hardened the world's randomness fabric and delivered open kits for others to verify. Either way, the reflexive architecture embodies Logostics: q(t) -> tau(t) not by assertion, but by transparent alignment—changing behavior when divergence rises, publishing proofs when it falls, and inviting the community to check every step.

References

Acín, A., & Masanes, L. (2016). Certified randomness in quantum physics. *Nature*, 540, 213–219.

Arnon-Friedman, R., Renner, R., & Vidick, T. (2018). Simple and tight device-independent security proofs by entropy accumulation. *Communications in Mathematical Physics*, 372, 123–165.

Bell, J. S. (1964). On the Einstein-Podolsky-Rosen paradox. *Physics*, 1, 195–200.

Bertoni, G., Daemen, J., Peeters, M., & Van Assche, G. (2013). *Keccak*. In *Advances in Cryptology – EUROCRYPT 2013* (pp. 313–314). (See also NIST FIPS 202: SHA-3, 2015.)

Brown, R. G. (2004). *Dieharder: A random number test suite*. Software report.

Carter, J. L., & Wegman, M. N. (1979). Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2), 143–154.

Clauser, J. F., Horne, M. A., Shimony, A., & Holt, R. A. (1969). Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15), 880–884.

Dupuis, F., Fawzi, O., & Renner, R. (2020). Entropy accumulation. *Communications in Mathematical Physics*, 379, 867–913. (Preprint 2016.)

Giustina, M., et al. (2015). Significant-loophole-free test of Bell's theorem with entangled photons. *Physical Review Letters*, 115, 250401.

Grieder, P. K. F. (2001). *Cosmic Rays at Earth: Researcher's Reference Manual and Data Book*. Elsevier.

Hensen, B., et al. (2015). Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526, 682–686.

Howard, S. R., Ramdas, A., McAuliffe, J., & Tibshirani, R. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2), 1055–1080.

Impagliazzo, R., Levin, L. A., & Luby, M. (1989). Pseudo-random generation from one-way functions. *Proceedings of the 21st ACM STOC*, 12–24. (Leftover hash lemma foundation.)

L'Ecuyer, P., & Simard, R. (2007). TestU01: A C library for empirical testing of random number generators. *ACM Transactions on Mathematical Software*, 33(4), 22.

Nickolaenko, A. P., & Hayakawa, M. (2014). *Schumann Resonances for Tyros*. Springer.

NIST. (2010). *SP 800-22 rev.1a: A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*.

NIST. (2015). *FIPS 202: SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions*.

NIST. (2018). *SP 800-90B: Recommendation for the Entropy Sources Used for Random Bit Generation*.

Pironio, S., et al. (2010). Random numbers certified by Bell's theorem. *Nature*, 464, 1021–1024.

Ramdas, A., Ruf, J., Larsson, M., & Koolen, W. M. (2023). Testing exchangeability: Fork-convexity, e-processes, and the probability integral transform. *The Annals of Statistics*, 51(1), 1–26. (E-values and anytime-valid testing foundations.)

Schumann, W. O. (1952). Über die strahlungslosen Eigenschwingungen einer leitenden Kugel, die von einer Luftschicht und einer Ionosphärenhülle umgeben ist. *Zeitschrift für Naturforschung A*, 7(2), 149–154.

Shafer, G. (2021). *The Language of Betting: A Probability and Finance Synthesis*. (E-variable/martingale perspective.)

Shalm, L. K., et al. (2015). Strong loophole-free test of local realism. *Physical Review Letters*, 115, 250402.

Ville, J. (1939). *Étude critique de la notion de collectif*. Gauthier-Villars. (Ville's inequality; anytime control.)

Willems, F. M. J., Shtarkov, Y. M., & Tjalkens, T. J. (1995). The context-tree weighting method: Basic properties. *IEEE Transactions on Information Theory*, 41(3), 653–664.

Ziv, J., & Lempel, A. (1977). A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3), 337–343.

Ziv, J., & Lempel, A. (1978). Compression of individual sequences via variable-rate coding. *IEEE Transactions on Information Theory*, 24(5), 530–536.

Standards and engineering practice (supporting).

- NIST SP 800-90A/B/C series (random bit generators; entropy sources; conditioning components).
- NIST SP 800-56C and SP 800-133 (extractor use in key establishment and key generation).
- Vendor datasheets for balanced photodiodes, APDs/SiPMs, TIAs, and high-speed ADCs (for specific BOM validation).

Appendices

A. Derivations for anytime-valid tests and martingale construction

A.1 Ville's inequality and e-process alarms

Let M_t be a nonnegative supermartingale under null P0 (iid Bernoulli(1/2)): E_0[M_t | F_{t-1}] <= M_{t-1}, M_0 = 1.
Ville's inequality: for any alpha in (0,1),
Pr_0(sup_{t>=1} M_t >= 1/alpha) <= alpha.
Thus the stopping rule "alarm at first t with M_t >= 1/alpha" has type-I error <= alpha with arbitrary optional stopping and peeking.

Define an anytime p-value p_t := 1 / max_{s<=t} M_s. Then Pr_0(inf_t p_t <= alpha) <= alpha.

A.2 Bias betting e-process for bits

Let X_t in {0,1}, Y_t := 2X_t−1 in {−1,+1}. Choose predictable stakes lambda_t in (−1,1). Define M_t = product_{i=1..t} (1 + lambda_i * Y_i).
Under P0 (E_0[Y_i | F_{i-1}] = 0), E_0[M_t | F_{t-1}] <= M_{t-1} by convexity, so M_t is a supermartingale.
Practical choice: lambda_t = clip(c * m_t, −lambda_max, +lambda_max), where m_t is a running mean of Y and c in (0,1). This yields power against persistent bias.

Mixture e-process (safe default). For x_t = sum X_i,
M_t = Beta(x_t + a, t − x_t + b) / Beta(a,b) * 2^{−t}.
This equals the Bayes factor of a Beta(a,b) alternative vs Bernoulli(1/2) and is an e-process for any fixed a,b>0.

A.3 Order-1 Markov leakage e-process

Estimate transition counts N_{u->v}(t) for u,v in {0,1}. Under iid(1/2), the likelihood per step is 1/2. Under the Markov alternative with MLE transition probs p_{u->v} = N_{u->v} / N_{u}, the likelihood ratio over t steps is
LR_t = product_{i=2..t} (p_{X_{i-1} -> X_i} / 1/2).
A Jeffreys or Dirichlet prior over transition rows yields a mixture LR that is an e-process (safe against overfitting). Report M_t = LR_t_mix.

A.4 Spectral sentinels as blockwise e-values

Split into blocks w of length B. For each block, compute targeted line energy T_w (e.g., squared magnitude of DFT at 50/60 Hz). Convert the one-shot null distribution to an e-value e_w = f(T_w) with E_0[e_w] <= 1 (calibrated via simulation or analytic bound). Then M_t = product_{w<=t} e_w is an e-process.

A.5 Universal compression e-process

Let L_t be cumulative codelength from an online universal code (CTW/LZ). For iid(1/2), expected codelength is t up to o(t). Define
M_t = 2^{ t − L_t }.
Kraft inequality and universality imply E_0[M_t] <= 1, hence an e-process. Persistent compressibility (L_t < t) makes M_t grow.

A.6 Anytime confidence sequences (CS) for mean

For theta = E[X], construct a mixture e-process over a grid of alternatives and invert to get CS_t(alpha) = { theta : max_{s<=t} e_s(theta) < 1/alpha }.
Practical: empirical-Bernstein CS with radius r_t = O(sqrt((v_hat_t * log log t + log(1/alpha))/t)), giving uniform-in-time coverage.

B. Entropy accumulation mapping: Bell score -> min-entropy rate

B.1 CHSH bookkeeping

Trials i=1..n with inputs (x_i,y_i) in {0,1}^2, outputs (a_i,b_i) in {0,1}^2. Define CHSH game wins W_i in {0,1}. Let w_hat = (1/n_pe) sum_{i in PE} W_i over parameter-estimation subset PE. Equivalently, S_hat = 8 w_hat − 4 in [2, 2*sqrt(2)].

Obtain a lower confidence bound S_lo (or w_lo) with failure prob delta_pe using concentration valid under memory (EAT-compatible bounds).

B.2 Per-trial entropy curve

Use a standard DI bound f_CHSH(S) giving min-entropy per A-outcome conditioned on inputs and quantum side info:
h_min_per_trial >= f_CHSH(S), with f_CHSH(2)=0, monotone increasing, f_CHSH(2*sqrt(2)) > 0. (Implement the specific analytic form from the chosen DI reference; document version in prereg.)

B.3 EAT aggregation and finite-size

Let n_prod be production trials. Entropy Accumulation Theorem yields
H_min^{epsilon_s}(A^{n_prod} | X^{n} E) >= n_prod * f_CHSH(S_lo) − penalty(n, q_pe, delta_pe, epsilon_s),
where penalty accounts for finite size, estimator slack, and smoothing. All constants, deltas, and the f_CHSH version are preregistered.

B.4 Extractable bits from DI lane

For extractor error epsilon_ext,
k_DI <= H_min^{epsilon_s} − safety_overhead(epsilon_ext) − leakage_budget.
safety_overhead for universal hashing is approx 2 * log2(1/epsilon_ext). Compose epsilons additively across windows and lanes.

C. Extractor parameters, target epsilon, and rate control math

C.1 Leftover hash lemma sizing (Toeplitz or XOF)

Given N raw bits with min-entropy H_min (smooth),
k <= H_min − 2*log2(1/epsilon_ext) − leakage_budget.
Choose epsilon_ext so that composed epsilon_total = epsilon_DI + epsilon_PH + epsilon_comp meets the system target.

Example (copy-safe).
N = 1.2e8, h_min = 0.62 per bit -> H_min = 7.44e7.
epsilon_ext = 2^(−128) -> overhead = 256 bits.
leakage_budget = 512 bits -> k <= 7.44e7 − 768.

C.2 Toeplitz matrix and seed handling

Toeplitz k x N specified by N+k−1 seed bits. Use a CSPRNG to expand a per-window master seed; commit seed_commit = H(seed). Never reuse seed across overlapping raw windows. Publish extractor_type, k, epsilon, seed_commit in the certificate.

C.3 Adaptive rate control

Let h_min_lo(t) be the conservative per-bit bound for the current window. Define target compression ratio rho(t) = k/N.

Policy:

- rho_req(t) = min_over_lanes(h_min_lo(t)) − overhead/N.
- If rho_req(t) < rho_current − Delta*, set rho_current := rho_req(t) (de-rate).
- If any monitor HOT or h_min_lo(t) falls below floor, set rho_current := 0 and halt.

Guarantee: with correct h_min_lo(t) and a 2-universal extractor, TV distance of output from uniform is <= epsilon_ext, independent of monitor outcomes.

C.4 Composition across windows and lanes

For sequential windows with epsilons {epsilon_i}, total distance <= sum epsilon_i (triangle inequality). Set per-window epsilon_i = epsilon_budget / W where W is sliding window count in the retention period. For XOR of DI and PH outputs, compose epsilons additively and take min of their entropy rates for sizing.

D. Preregistration checklist, locked-analysis scripts, and telemetry schema

D.1 Preregistration checklist (succinct)

- Hypotheses: primary endpoint (Delta*, direction), secondary endpoints, null regions.
- Design: spacelike separation L, latency budgets, shielding spec, duty cycles (DI vs PH).
- Alarms: alpha_hard, alpha_soft, spending schedule, HOT/WARM thresholds.
- Analysis: exact binaries and hashes; configs for monitors; block sizes; CS flavor and delta.
- Signaling protocols: block designs for human and automated schedules; sham and swap days; unblinding rules.
- Replication plan: number of independent days, site matrix, operator swaps.
- Data policy: release tiers, retention times, redaction rules.
- Loopbacks: tickle amplitudes, frequencies, and detection-power claims.
- Environmental regressors: list and preprocessing; residualization recipe.
- EAT specifics: q_pe, f_CHSH reference, penalties, epsilon_s.

D.2 Locked-analysis script outline (pseudo)

```
main():
  cfg = load_locked_config()          # hashed prereg config
  ensure_binary_hash(cfg.expected_hash)
  rng = init_csprng(cfg.seed_analytics)
  while acquire:
    W = read_raw_window()
    T = read_telemetry_window()
```

```
  h_min = bound_entropy(T, W, cfg.estimators_90B)
  k, eps = size_extractor(h_min, cfg.epsilon_target, cfg.leakage_budget)
  stats = run_monitors(W, cfg.monitor_set, cfg.spending)
  rho = policy_rate_control(h_min, stats, cfg.policy)
  cert = build_certificate(W.meta, k, eps, h_min, stats, T.digest, cfg)
  sign_and_append(cert)
  if stats.state == HOT or rho == 0:
    quarantine_and_halt()
```

D.3 Telemetry schema (JSON exemplar)

```
{
 "window_id": "A-2025-10-30T15:02:18Z-000173",
 "calib_version": "v1.3.2",
 "sensors": {
  "temp_board_C": {"mean": 33.4, "std": 0.2},
  "temp_optics_C": {"mean": 29.1, "std": 0.1},
  "vibration_rms_mg": 1.8,
  "psu_12V_ripple_mVrms": 3.2,
  "rf_probe_dBm": -72.5,
  "lo_power_mW": 4.6,
  "adc_hist": {"bins": [..], "counts": [..]},
  "dark_rate_kcps": 12.4
 },
 "bounds_inputs": {
  "afterpulse_prob": {"estimate": 0.003, "ci95": [0.002,0.004]},
  "dead_time_ns": {"estimate": 45, "ci95": [40,50]}
 }
```

}

Digest this object (canonical CBOR) and place the hash in telemetry_digest inside the certificate.

E. Human-intent variant: block designs, blinding, debrief, and safeguards

E.1 Block designs

- **FOCUS vs REST blocks:** durations 2–5 min each; K pairs per session; randomize order with constraints (no >2 consecutive identical blocks).
- **Direction precommit:** define effect direction before run (e.g., bias toward 1s during FOCUS). No mid-run changes.
- **Jittered edges:** add 5–15 s random jitter to block starts to reduce edge artifacts while keeping labels precise at B.

E.2 Blinding and operator protocol

- A receives only opaque block_id tokens. Dashboards never show label-conditioned summaries.
- B seals schedule with signature and time-stamps. No communication A<->B during acquisition.
- Environmental covariates recorded but withheld from A until unblinding.

E.3 Debrief and reporting

- After unblinding, report primary endpoint Delta_hat with anytime-valid p_value, confidence sequence, and ablations (trim edges, time-shifts, sham).
- Include a FOCUS self-report form (arousal, fatigue) strictly *post hoc* and clearly separated from statistical analysis.
- Publish nulls and positives identically: same plots, same tables, same prereg thresholds.

E.4 Safeguards against expectancy and leakage

- **Silence and light discipline:** no speech, displays, or LEDs visible between rooms; headphones optional white noise.
- **RF discipline:** phones off and out; spectrum monitor logs; Faraday bags for laptops if present.

- **Automation days:** alternate human days with automated schedules to test whether any effect depends on intent.

E.5 Ethics and well-being

- Participation is voluntary, with rest breaks; no coercion or performance demands.
- No invasive sensors. Optional HRV or EEG only under separate IRB and prereg.
- Clear language: the experiment does not test individual “powers”; it tests operational steerability under strict controls.

Net effect. The human-intent variant is fully nested within the same preregistered, blinded, and anytime-valid framework. If an effect appears, the next test is whether it survives replacement by automated schedules; if it does not, subsequent studies can isolate state-dependent correlates under the same gold-standard discipline.